



US 20030195877A1

(19) **United States**(12) **Patent Application Publication****Ford et al.**(10) **Pub. No.: US 2003/0195877 A1**(43) **Pub. Date: Oct. 16, 2003**

(54) **SEARCH QUERY PROCESSING TO
PROVIDE CATEGORY-RANKED
PRESENTATION OF SEARCH RESULTS**

(60) Provisional application No. 60/169,570, filed on Dec. 8, 1999.

Publication Classification

(51) Int. Cl.⁷ **G06F 7/00**

(52) U.S. Cl. **707/3**

(76) Inventors: **James L. Ford**, Bellevue, WA (US);
Christopher L. Scofield, Seattle, WA
(US); **Dwayne E. Bowman**,
Woodinville, WA (US); **Ruben E.
Ortega**, Seattle, WA (US)

Correspondence Address:

KNOBBE MARTENS OLSON & BEAR LLP
2040 MAIN STREET
FOURTEENTH FLOOR
IRVINE, CA 92614 (US)

(21) Appl. No.: **10/404,992**

(22) Filed: **Mar. 31, 2003**

Related U.S. Application Data

(63) Continuation of application No. 09/528,127, filed on Mar. 17, 2000.

(57) ABSTRACT

A search engine system displays the results of a multiple-category search according to levels of relevance of the categories to a user's search query. A query server receives a search query from a user and identifies, within each of multiple item categories, a set of items that satisfy the query. The sets of items are used to generate, for each of the multiple categories, a score that reflects a level significance or relevance of the category to the search. The scores may be based, for example, on the number of hits within each category relative to the total number of items in that category, the popularity levels of items that satisfy the query, a personal profile of the user, or a combination thereof. The categories are then presented to the user, together with the most relevant items within each category, in the order of highest to lowest category relevance.

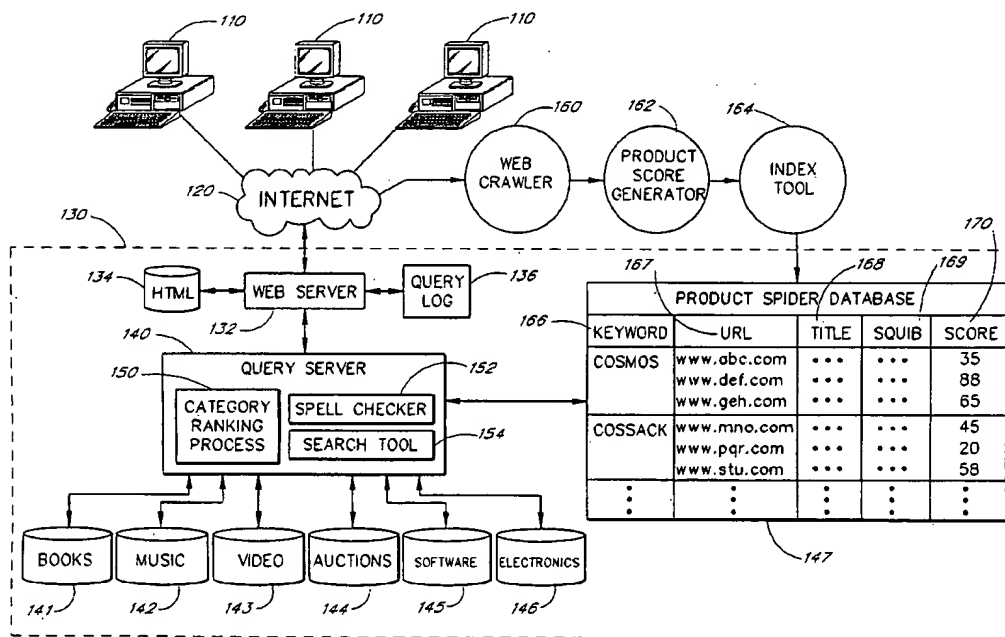



FIG. 1



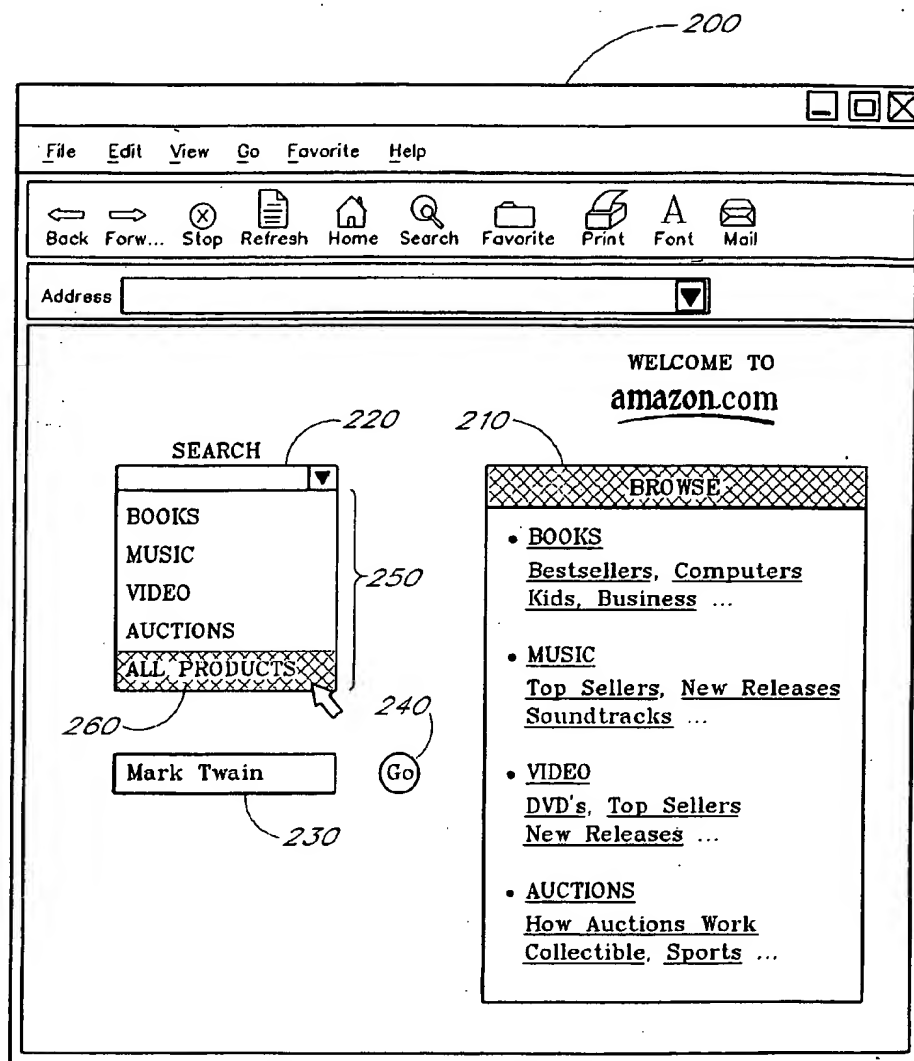


FIG. 2

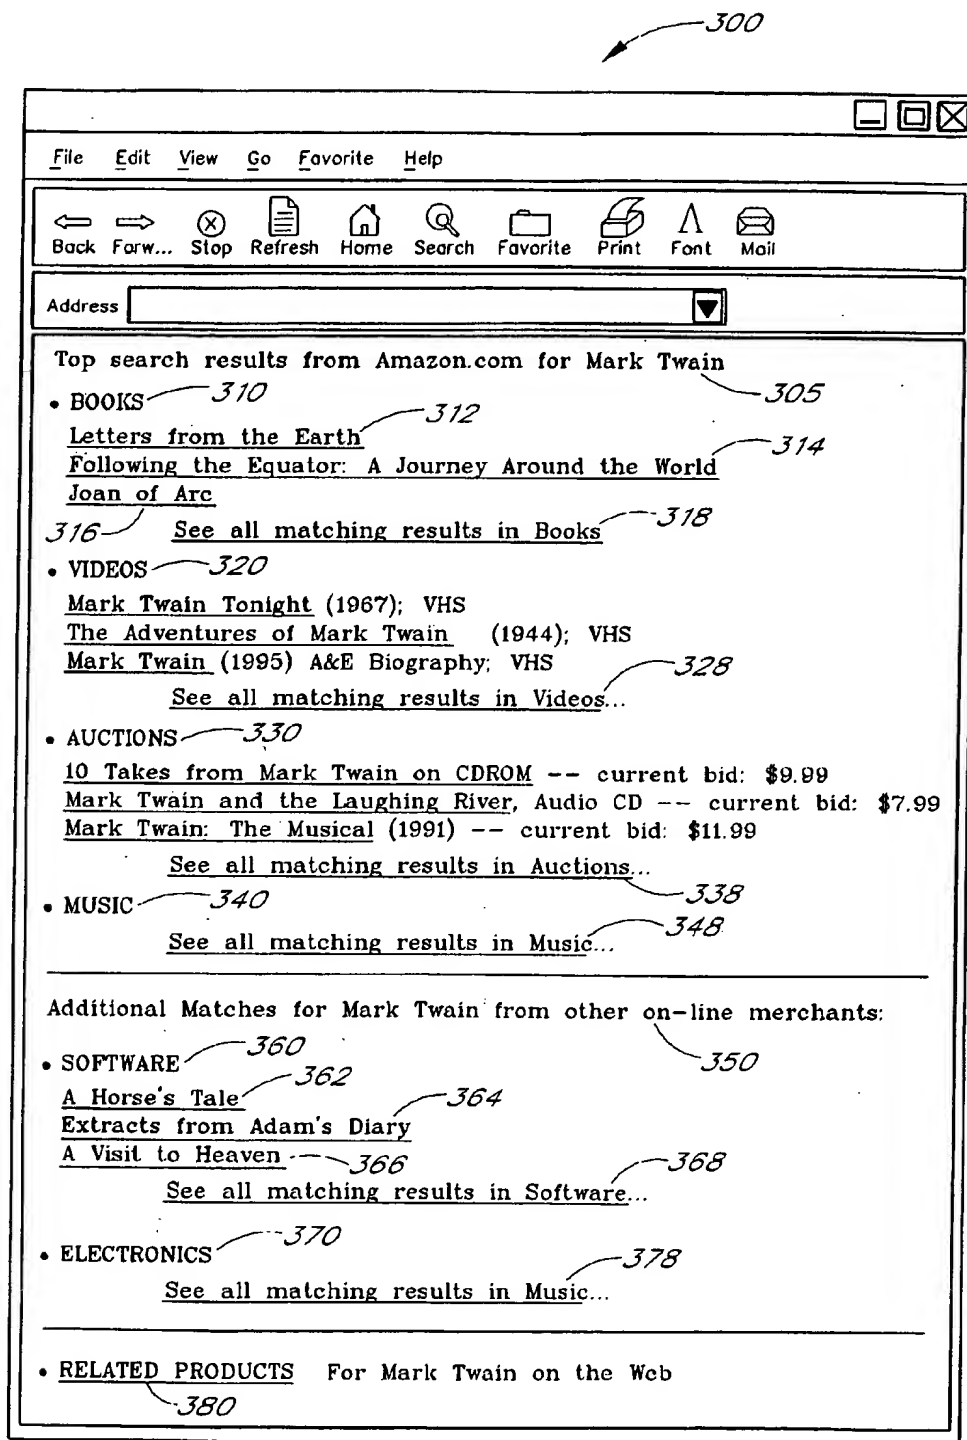


FIG. 3

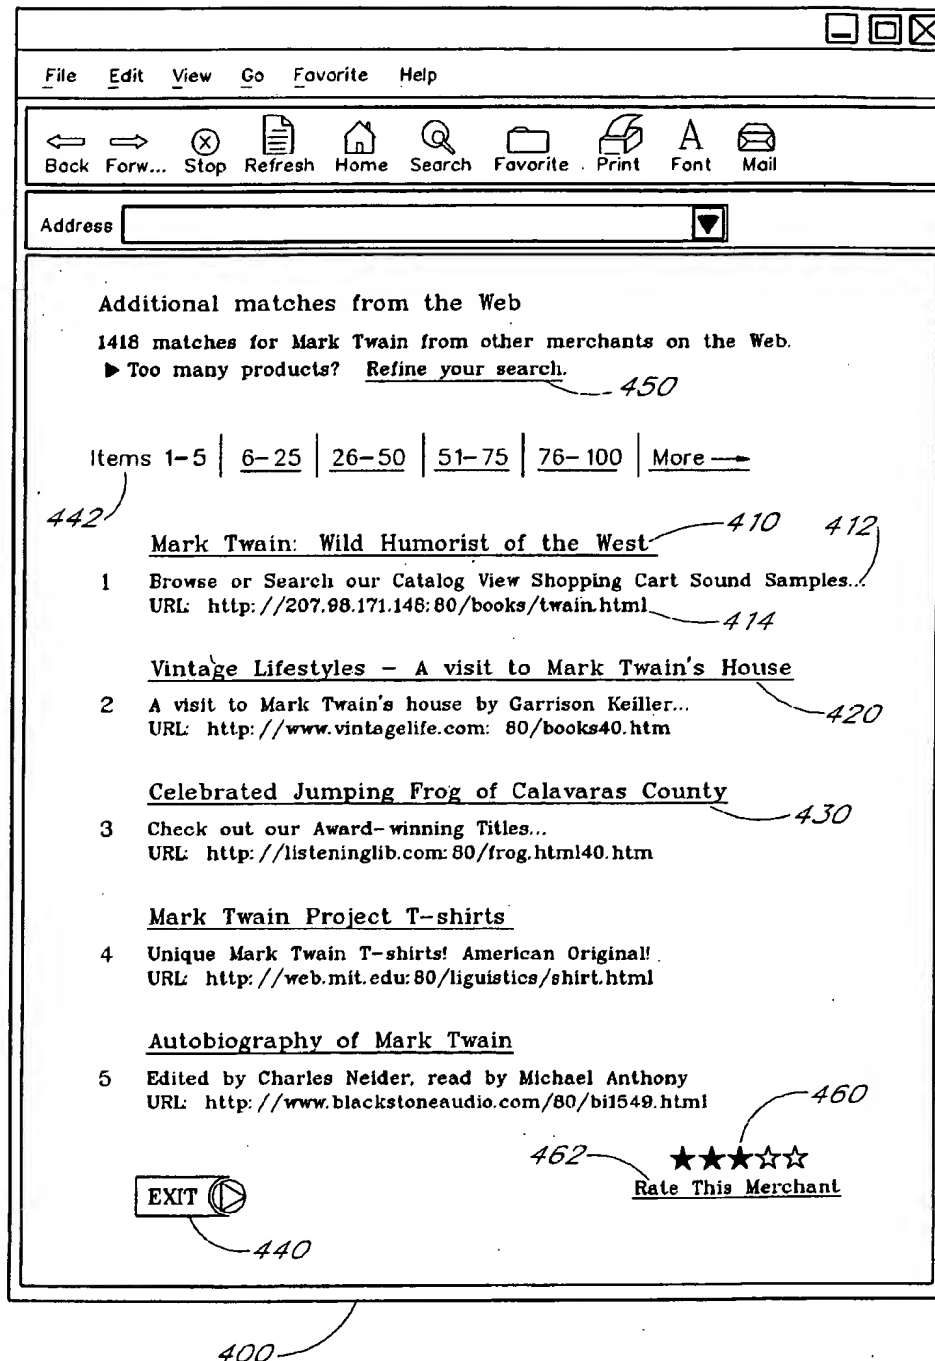


FIG. 4

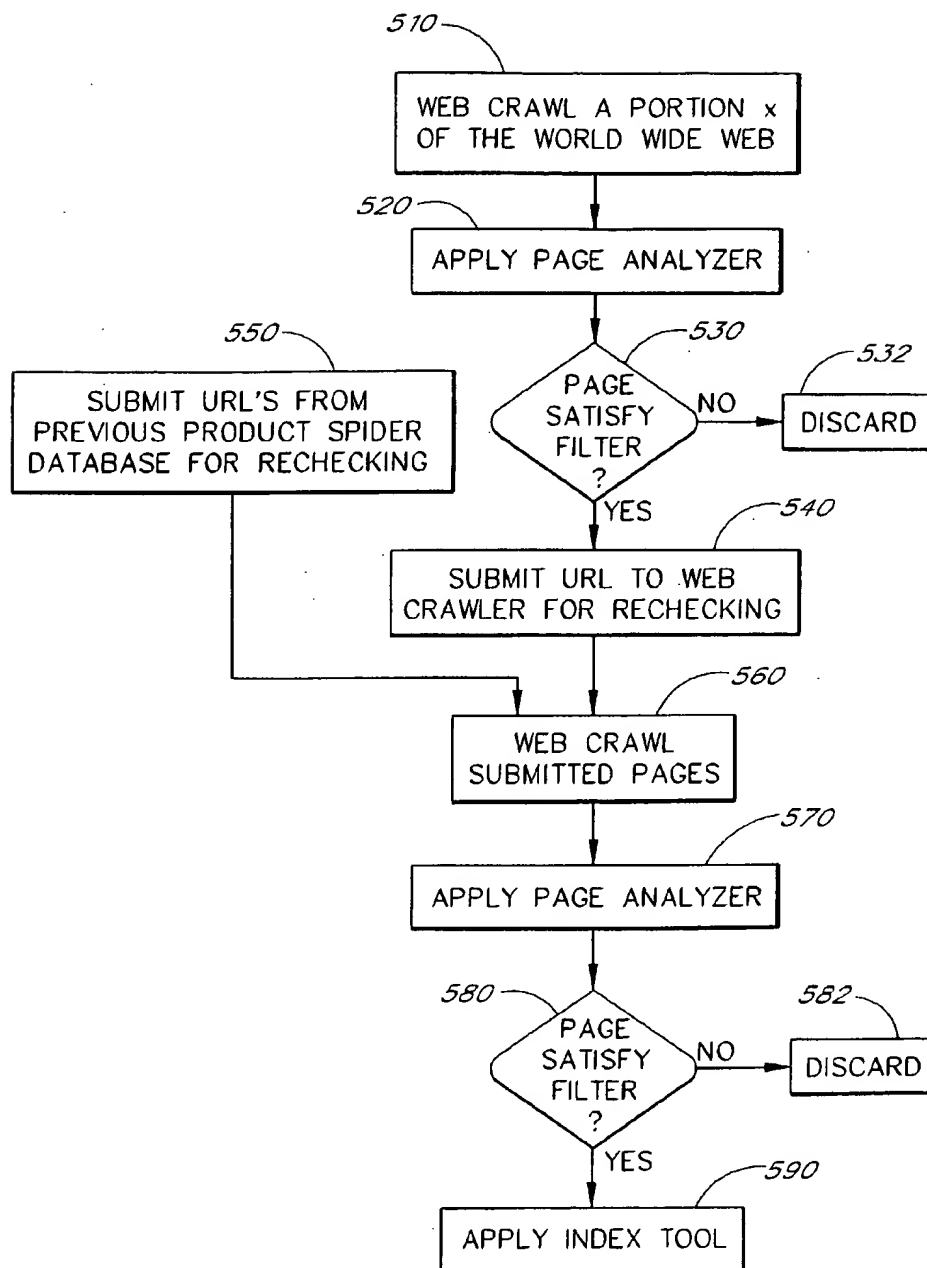


FIG. 5

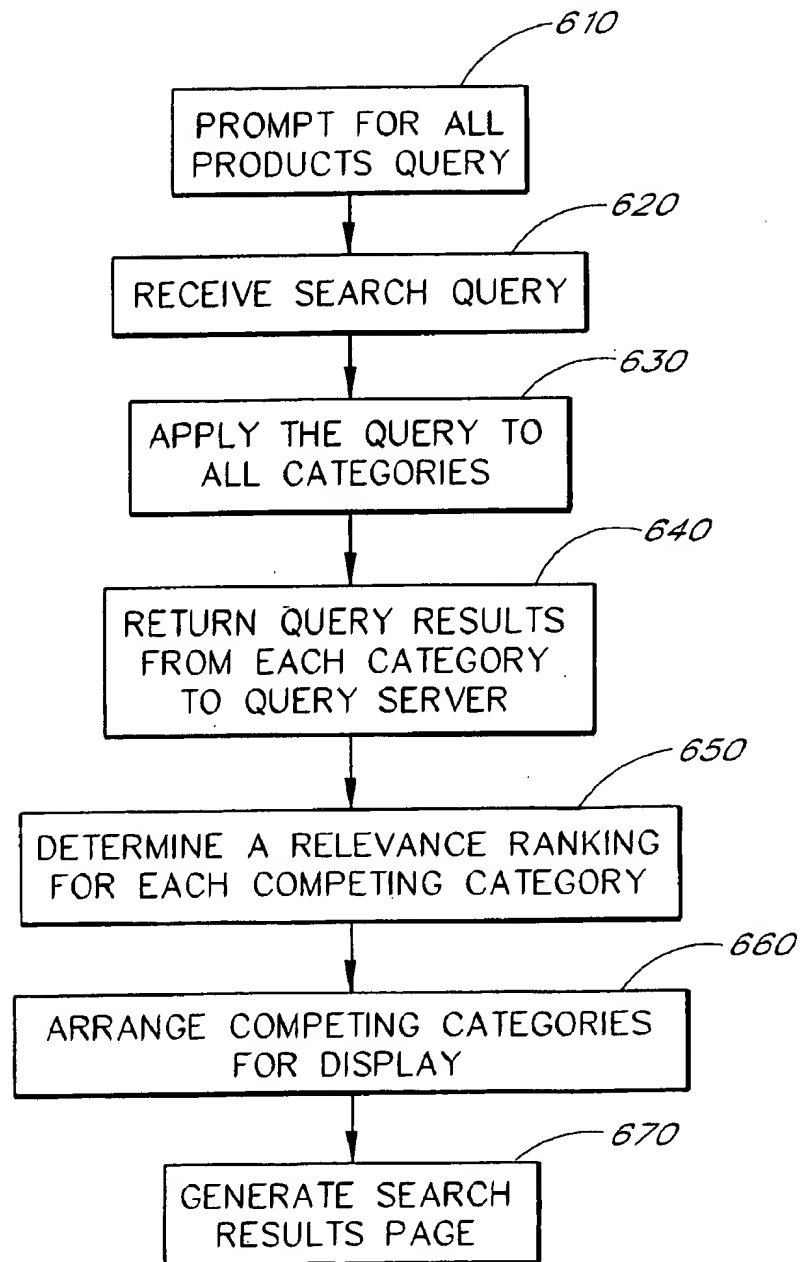


FIG. 6

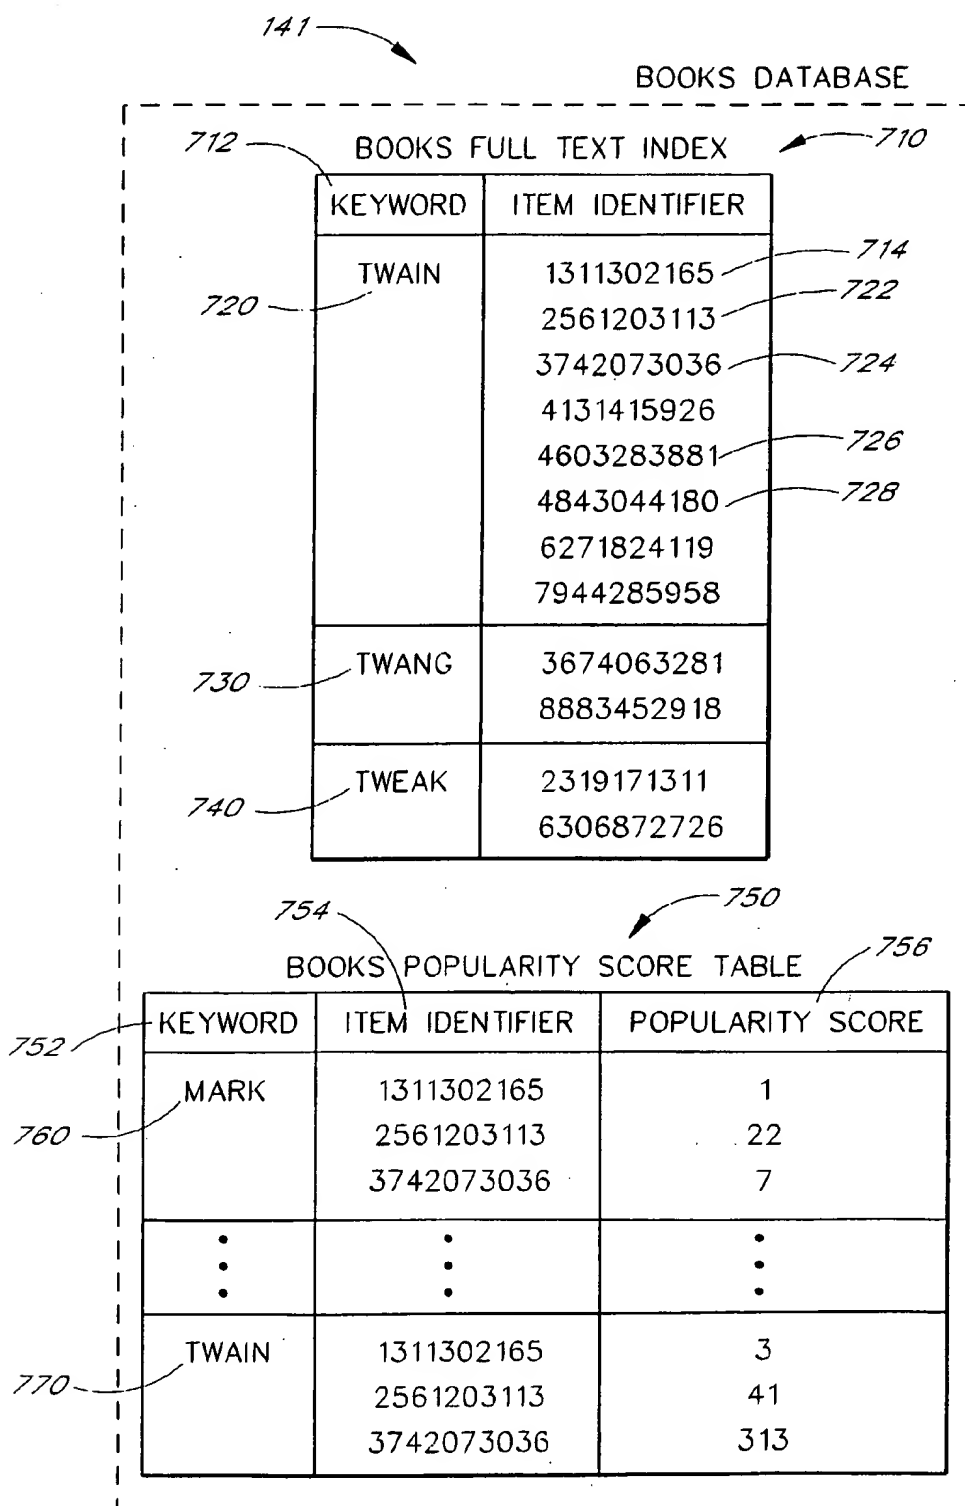


FIG. 7

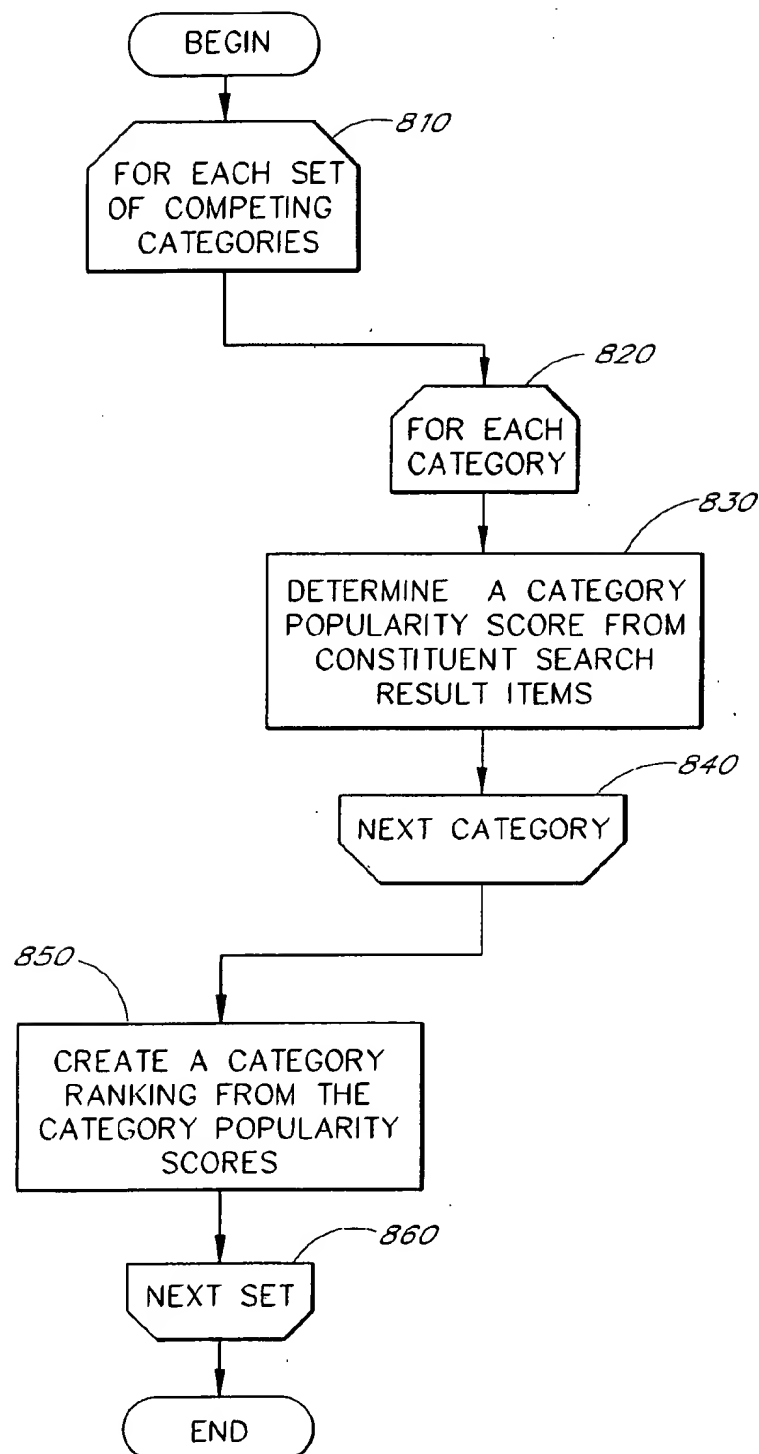
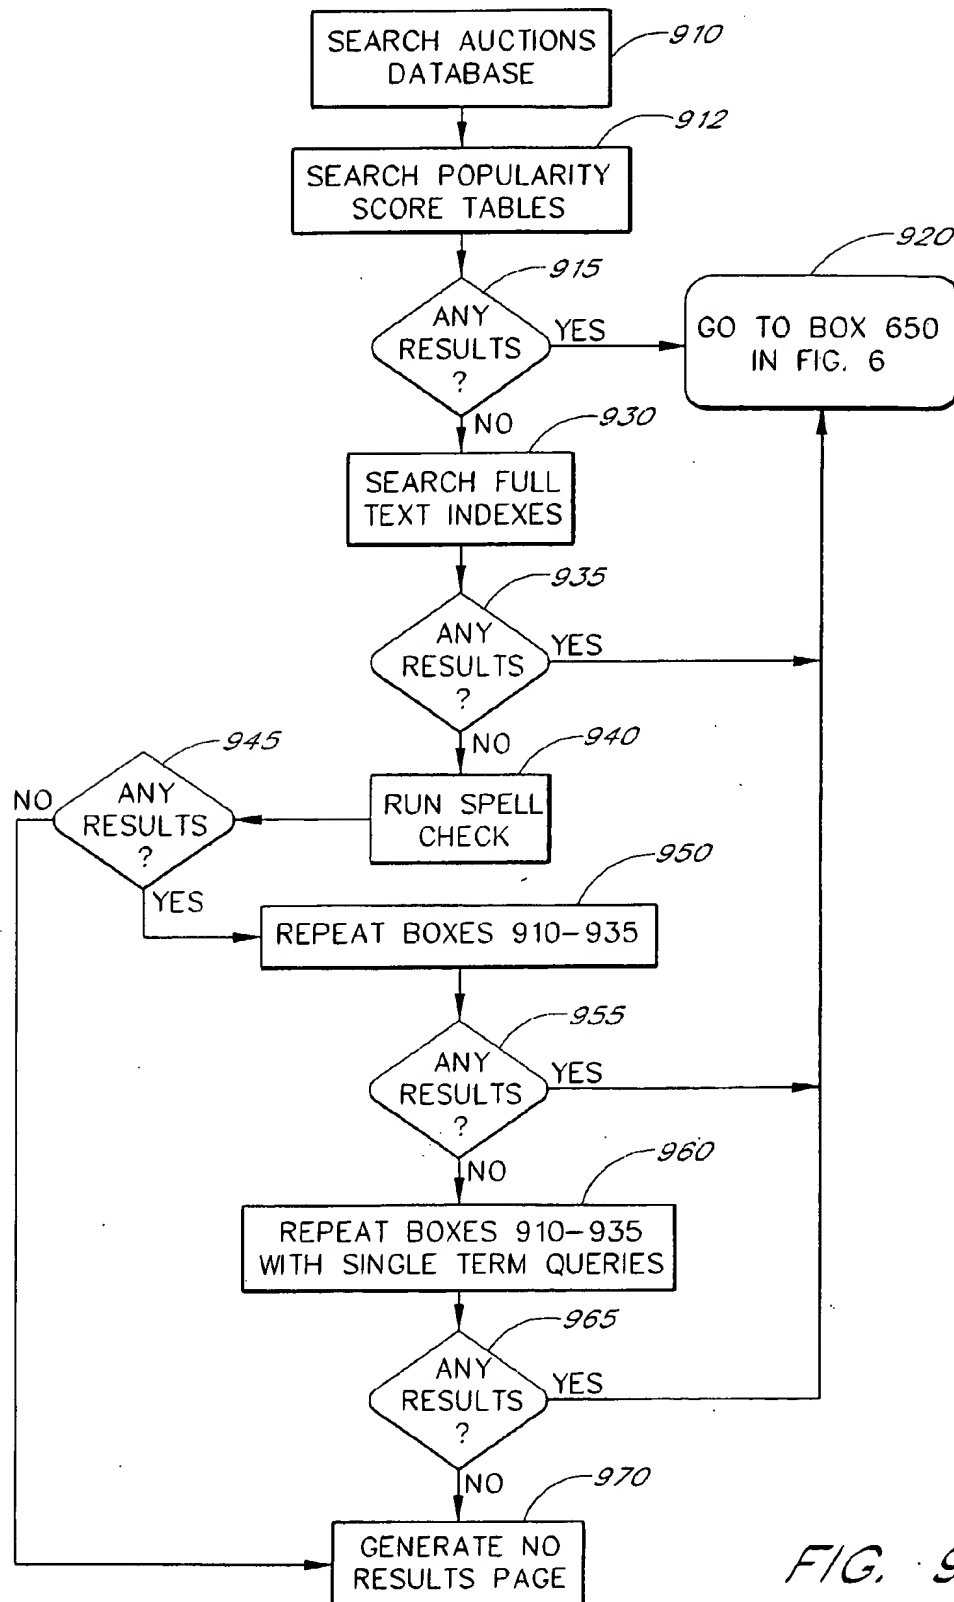


FIG. 8



SEARCH QUERY PROCESSING TO PROVIDE CATEGORY-RANKED PRESENTATION OF SEARCH RESULTS

RELATED APPLICATIONS

[0001] This application is a continuation of U.S. application Ser. No. 09/528,127, filed Mar. 17, 2000, the disclosure of which is hereby incorporated by reference.

FIELD OF THE INVENTION

[0002] The present invention relates to the field of search engines. More specifically, the invention relates to techniques for facilitating viewing search results that span multiple item categories, and for locating web pages that include offerings for products and other types of items.

BACKGROUND OF THE INVENTION

[0003] In the field of electronic commerce, it is common for online merchants to sell products within many different product-related categories. For example, Amazon.com, Inc., the assignee of the present application, sells products within the categories of books, music, video & DVD, toys & games, electronics, home improvements and auctions. The predefined categories and associated products are typically presented to users in the form of a browse tree. In addition, many merchants provide a search engine for conducting searches for products.

[0004] One problem commonly encountered by online merchants is the inability to effectively present groups of related products that span the predefined categories. Due to the large number of products and categories, and the organization of the web site, many relationships between products may be difficult for the user to ascertain. For example, suppose a user of a merchant's web site is a fan of the American humorist and author Mark Twain. The user may choose to look for books written by Mark Twain through a browse tree in the book section of an on-line commerce web site. Browsing in this manner is likely to reveal a large number of books authored by, or written about, Mark Twain. The user, however, may be unaware that the web site also sells products other than books that may be of interest to fans of Mark Twain. For example, a videos section of the same web site may contain video biographies of Mark Twain and video adaptations of many of his classic books, while a music section may include compact discs with songs inspired by his writings. Similarly, an auctions section of the same site may contain products offered for sale by third parties that may be of interest to the user, such as Mark Twain memorabilia. Although use of the web site's search engine may reveal some of these additional products, the user typically must review a long list of search results in order to identify the products or categories of interest.

[0005] Another problem in the field of on-line commerce is that of locating a web site from which a particular product can be purchased. This problem may arise, for example, when the online merchants known to the consumer do not carry the product of interest. In such a circumstance the consumer may use an Internet search engine such as ALTAVISTA or EXCITE to search for a web site that sells the product. The scope of such a general search is often large enough, however, that only a small fraction of a large number of located web sites actually offer the product for

sale. For example, the search may include a relatively large number of sites that merely provide reviews, technical support, specifications, or other information about the product of interest. Thus the sites of greatest interest to the consumer are likely to be buried deep within a long list.

[0006] The present invention seeks to overcome these and other problems.

SUMMARY OF THE INVENTION

[0007] The present invention provides various features for assisting users in conducting online searches. The features may be embodied alone or in combination within a search engine of an online merchant, an Internet search engine, or another type of search system.

[0008] One feature of the invention involves a method for displaying the results of a multiple-category search according to levels of significance of the categories to a user's search query. The method can be used to display the results of a search for products or for any other type of item. In a preferred embodiment, the method involves receiving a search query from a user and identifying, within each of multiple item categories, a set of items that satisfy the query. The sets of items are then used to generate, for each of the multiple categories, a score that indicates a level significance or relevance of the category to the search. The scores may be based, for example, on the number of hits (items satisfying the query) within each category relative to the total number of items in that category, the popularity levels of items that satisfy the query, or a combination thereof.

[0009] The categories and associated items are then presented to the user in a display order that depends upon the scores—preferably from highest-to-lowest significance. Other significance criteria, such as a category preference profile of the user, may additionally be used to select the display order. In addition, other display methods for highlighting the most highly ranked categories may additionally or alternatively be used. The method increases the likelihood that the categories that are of most interest to the user will be presented near the top of the search results listing, or otherwise called to the attention of the user. To assist the user in efficiently viewing a cross section of the located items and their categories, no more than N items (e.g., the most highly ranked three items) within each category are preferably displayed on the initial search results page.

[0010] Another feature of the invention involves a system and methods for assisting users in locating web sites or pages from which user-specified products can be purchased. In a preferred embodiment, each web page located by a crawler program is initially evaluated, according to a set of content-based rules, to generate a score that indicates a likelihood that the web page includes a product offering. The scores may additionally be based on other criteria, such as the content of other web pages of the same web site. Representations of some or all of the scored web pages are stored in a keyword index that maps keywords to addresses (URLs) of the web pages. The keyword index is used by a query server to locate web pages that are both relevant to a user's search query and likely to include a product offering. This may be accomplished, for example, by limiting a scope of the search to web pages having a score that satisfies a particular threshold.

[0011] In one embodiment, the above-described features are embodied in combination within a search engine of a host merchant's web site. From this web site, a user can initiate an "All Products" type search that spans multiple product categories. The submitted search query is used to identify a set of products that satisfy the query, and a set of web pages that both satisfy the query and that have been determined to likely include product offerings. The results of the search are presented using a composite web page which lists at least some of the located products and at least some of the located web pages. The products are preferably displayed in conjunction with their respective product categories according to the above-described category ranking and display method.

BRIEF DESCRIPTION OF THE FIGURES

[0012] These and other features and advantages of the invention will now be described with reference to the drawings of certain preferred embodiments, which are intended to illustrate and not to limit the invention, and in which:

[0013] FIG. 1 illustrates a system in which users access web site information via the Internet, and illustrates the basic web site components used to implement a search engine that operates in accordance with the present invention.

[0014] FIG. 2 illustrates a sample search tool interface page of the web site.

[0015] FIG. 3 illustrates a sample results page for an All Products search. The results include items directly offered for sale by the host merchant web site, items offered for sale by third parties using the host web site as a forum, items offered for sale by on-line merchants affiliated with the host merchant, and items offered for sale by on-line merchants unaffiliated with the host merchant.

[0016] FIG. 4 illustrates a sample results page displaying "Related Products" items associated with on-line merchants who are unaffiliated with the host merchant web site.

[0017] FIG. 5 illustrates the process used to generate the product spider database of FIG. 1.

[0018] FIG. 6 illustrates the process used to generate a return page in response to an "All Products" search query.

[0019] FIG. 7 illustrates the structure of the Books database of FIG. 1.

[0020] FIG. 8 illustrates the process used to generate a category relevancy ranking for use by the process of FIG. 6.

[0021] FIG. 9 illustrates the process used to find All Products search results for both common and uncommon search queries.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0022] One feature of the present invention involves a method for identifying and displaying product information derived from multiple product categories to a user in response to a query submitted to a search engine by the user. Another feature of the invention involves methods for users to search for web pages from which particular products can be purchased. In the preferred embodiment, these two features are embodied within a common search engine system;

as will be apparent, these and other features of the invention can be used independently of one another and may therefore be considered as distinct inventions. For convenience of description, however, the term "invention" is used herein to refer collectively to the various inventive features disclosed.

[0023] A preferred embodiment and implementation of the invention will now be described with reference to the drawings. The description will reference various details of the invention in the context of AMAZON.COM's web site. These details are set forth in order to illustrate, and not to limit, the invention. The scope of the invention is defined only by the appended claims.

[0024] A. Overview of Web Site and Search Engine

[0025] FIG. 1 illustrates the AMAZON.COM web site 130, including components used to implement a search engine in accordance with the invention. As is well known in the art of Internet commerce, the AMAZON.COM web site includes functionality for allowing users to search, browse, and make purchases from an on-line catalog of book titles, music titles, and other types of items via the Internet 120. Because the catalog contains millions of items, it is important that the site provide an efficient mechanism for assisting users in locating items.

[0026] As shown in FIG. 1, the web site 130 includes a web server application 132 ("web server") that processes user requests received from user computers 110 via the Internet 120. These requests include queries submitted by users to search the on-line catalog for products. The web server 132 records the user transactions, including query submissions, within a query log 136.

[0027] The web site also includes a query server 140 that processes queries by searching a number of databases 141-147. The Books database 141, Music database 142, and Videos database 143, include product identifiers for books, musical products, and multimedia products, respectively, that users may purchase directly from the web site 130. The AMAZON.COM web site includes other categories of products sold directly through the web site, such as Electronics and Toys & Games, that are omitted from FIG. 1 in the interest of clarity. The Books, Music, and Videos databases 141-143 are intended to represent all databases within the web site 130 associated with products marketed directly by the web site merchant.

[0028] The Auctions database 144 of FIG. 1 includes information about third party on-line auctions hosted by the web site 130. The AMAZON.COM web site also hosts fixed-price third party offerings, known as "zShops," corresponding to a version of an on-line "flea market." The zShops category contains a database analogous to the Auctions database 144 that is omitted from FIG. 1 in the interest of clarity. The Auctions database 144 in FIG. 1 is intended to represent all databases within the web site 130 associated with hosting third party transactions.

[0029] The Affiliated Merchant databases labeled Software 145 and Electronics 146 include information about software and electronics products, respectively, that are offered for sale on independent web sites affiliated with the host web site 130. The AMAZON.COM web site includes other categories of products sold on independent affiliated web sites, such as Sports & Outdoors and Toys & Games, that are omitted from FIG. 1 in the interest of clarity. The

Software and Electronics databases **145**, **146** are intended to represent all databases associated with products sold by independent web site merchants affiliated with the web site **130**.

[0030] The Product Spider database **147** includes information about independent web sites, unaffiliated with the host web site **130**, that have been identified as offering products for sale. This database is particularly useful in that it allows the host web-site **130** to help a consumer find product offerings for products that are not sold by the host web site **130** or by affiliated on-line merchants.

[0031] Each of the databases **141-147** contain data tables indexed by keyword to facilitate searching in response to queries. For simplicity, the division of the product offerings into multiple databases will be referred to as separate "categories." For example, the Product Spider database **147** is one of seven categories displayed in FIG. 1.

[0032] The web site **130** also includes a database of HTML (Hypertext Markup Language) content that includes, among other things, product information pages that show and describe products associated with the web site **130**.

[0033] The query server **140** includes a category ranking process **150** that prioritizes, by category, the results of searches across all of the various databases **141-147**. The prioritization scheme is based upon an assessment of the significance of each category to the search query submitted by the user. The query server **140** also includes a spell checker **152** for detecting and correcting misspellings in search attempts, and a search tool **154** capable of generating search results from a database (e.g. the Books database **141**) in response to a query submitted by a user. The search tool **154** prioritizes the items within a search result using different criteria depending upon the database used for the search. One approach, used for the Product Spider database **147**, ranks the search result items through the well known "term frequency inverse document frequency" (TFIDF) approach, in which the weighting applied to each term of a multiple-term query is inversely related to the term's frequency of appearance in the database. In other words, the term in a query that appears least often in a database (e.g. the Product Spider database **147**) is considered to be the most discriminating term in the query, and thus is given the greatest weight by the search tool **154**. Algorithms for implementing this approach are well known and are commonly available in software development kits associated with commercial search engines such ALTAVISTA and EXCITE.

[0034] The Product Spider database **147** is generated through the use of a web crawler **160** that crawls web sites on the Internet **120** while storing copies of located web pages. The output of the web crawler **160** is input to a product score generator **162** that assigns a numerical score ("product score") to each web page based upon the likelihood that the page offers a product for sale for either online or offline purchase. For purposes of generating the score in the preferred embodiment, any type of item that can be purchased is considered a "product," including but not limited to physical goods, services, software, and downloadable content. In other embodiments, the products may be based on a more narrow definition of what constitutes a product. For example, by requiring or taking into account whether a web site includes information about shipping, non-physical items can be excluded from consideration or

accorded a lesser weight. As depicted in FIG. 1, the product score **170** associated with each indexed web page is stored in the Product Spider Database **147**. Alternatively, the web page entries could be grouped according to product score (e.g., top third, middle third, bottom third) without actually storing the score values. As a further refinement, the product scores could be generated and stored on a site-by-site basis rather than on a page-by-page basis.

[0035] Information within the web pages assessed by the product score generator **162** is extracted by an index tool **164** and stored into the Product Spider database **147**. In some embodiments, only those web pages for which the product score exceeds a threshold value are indexed by the index tool **164**. The index tool **164** is complimentary to the search tool **154** in that the index tool **164** outputs data in a fully text indexed database format that is searchable by the search tool **154**. Any of a variety of commercially available search engine products, such as those available from ALTAVISTA and EXCITE, may be used to implement the index and search tools **164**, **154**.

[0036] As a variation to the above-described method, the index tool **164** could be configured to extract only those keywords that fall within a predefined distance (e.g., 10 words) of indicia of a product offering. This distance can be a fixed distance, or can be selected based on the type of indicia involved (dollar sign, manufacturer name, etc). This variation would tend to produce a keyword index in which the keywords are associated with specific products. A similar approach can be used to generate the squib **169**; for example, the squib could be generated by extracting sentences that immediately precede or follow some indicia of a product offering.

[0037] As noted above, the Product Spider database **147** is indexed by keyword **166**. Each keyword in the database is associated with one or more web pages for which the indexer **164** has determined an association. For each keyword-web site combination, the database includes a URL ("uniform resource locator") address **168** for locating the web site, a short string of text (a "squib") **169** extracted from the web site, and a product score value **170** indicative of the likelihood that the web site offers a product for sale.

[0038] The Product Spider database **147** may include information beyond that shown in FIG. 1. Other types of information that may be stored include, for example, a product category (e.g., books, music, video, etc.) ascertained from parsing the web page (or a collection of pages), an age appropriateness indicator (e.g., products appropriate for adults only), a language indicator for the web site (English, Spanish, etc.), product reviews, whether the offers are for new or used products, and whether the products are available on-line, off-line, or both. Furthermore, in embodiments in which the web site **130** provides users an option to rate the located merchants (e.g., on a scale of 1-5), and to view the ratings entered by other users, the Product Spider Database **147** may store the merchant ratings data.

[0039] The web server **132**, query server **140**, category ranking process **150**, and database software run on one or more UNIX-based servers and workstations (not shown) of the web site **130**, although other platforms could be used. To accommodate a large number of users, the query server **140** and the databases **141-147** may be replicated across multiple machines. The web site components invoked during the

searching process are collectively referred to herein as a "search engine." The web crawler 160 is preferably running continuously on one or more platforms (not shown) separate from the platforms used for the search engine. The product score generator 162 and indexer 164 preferably run on one or more platforms (not shown) separate from those used for the search engine and web crawler 160.

[0040] FIG. 2 illustrates the general format of a search tool interface page 200 of the host web site 130 that can be used to search for products. Users can pursue products using a browse tree interface 210 organized into predetermined categories such as books, music, videos, and auctions. Alternatively, users may search for products using a search engine interface 220. Users can perform searches with the search engine interface 220 by typing in the desired information (referred to herein as a "query") into a query window 230 and then clicking on a search initiation button 240. The user may control the scope of the search with a pulldown window 250 containing multiple categories. The search may be limited to any one category through selection of that category from the pulldown menu 250. Alternatively, the user may conduct a broad-based search through selection of an "All Products" option 260.

[0041] If the query is submitted to a single category, the search engine will present to the user a query results page (or multiple pages linked by hypertext, if the search finds a large number of items) containing a list of items matching the query. The search results page includes, for each item found, a hypertext link to additional web pages containing, among other things, product information about the item.

[0042] For multiple-term queries, the query server 140 effectively logically ANDs the query terms together to perform the search. For example, if the user enters the terms "Mark" and "Twain" into the query field 230 of FIG. 2, the query server 140 will search for and return a list of all items that are associated with both words.

[0043] If the search fails to find a single matching item, the search engine seeks to find misspellings within the query by submitting each query term to the spell checker 152. Any of a variety of types of spell checkers may be used for this purpose. In one embodiment, the spell checker 152 operates as described in U.S. application Ser. No. 09/115,662, filed Jul. 15, 1998, entitled "System and Method for Correcting Spelling Errors in Search Queries," which is hereby incorporated by reference. If the spell checker 152 determines that a term of the search query may be misspelled, a new term is substituted into the query and the query server 140 completes a new search with the modified query. In this situation the user is notified of the modification made to the query. If no results are found with the modified query, the user is presented with a "no results" page.

[0044] If no results are found that contain all of the query terms in a multiple-term query, the query server 140 performs searches on the individual terms of the query. In this situation the user is notified of the absence of exact matches, and is informed that the results merely represent close matches.

[0045] A more complete discussion of the processing of misspelled, or otherwise unusual, query terms is reserved until later with the help of FIG. 9.

[0046] When the user submits a query from the search engine interface 220 of FIG. 2 to the web site 130, the query

server 140 applies the query to the database, or databases, corresponding to the search scope selected by the user. For example, if the user has selected the "Books" field from the pulldown menu 250, the query is only submitted to the Books database 141.

[0047] If the user has selected the "All Products" field 260 of FIG. 2 from the pulldown menu 250, the query server 140 applies the query to all of the product databases 141-147. The search tool 154 generates results independently from each of the databases 141-147.

[0048] FIG. 3 illustrates the general format of a search results page 300 of the AMAZON.COM web site 130 generated and displayed to the user in response to an "All Products" search on the query "Mark Twain." The results page 300 displays the search results in three separate sections: a "Top Search Results" section 305, an "Additional Matches" section 350, and a "Related Products" section 380.

[0049] As shown in FIG. 3, the most prominent (i.e. highest) section of the page displays the Top Search Results section 305. This section displays some result items generated from application of the query to the databases directly associated with the web site 130, that is, to the Books, Music, Videos, and Auctions databases 141-144. These results are depicted in FIG. 3 categorized under the Books 310, Videos 320, Auctions 330, and Music 340 headings, and represent products that are available for purchase from the host merchant. For these categories, as many as three items associated with each category are preferably displayed to the user. The matching search result items displayed on the All Products search results page 300 are referred to as "top-level" search result items. For example, the top-level items listed under Books 310 are entitled "Letters From the Earth" 312, "Following the Equator: A Journey Around the World" 314, and "Joan of Arc" 316. Each top-level listing includes a hypertext link to product detail pages including information about the associated item. Hypertext links (318, 328, 338, 348) providing access to lists of additional items within the respective categories that match the search query are provided for each category as well. These non-displayed additional items are referred to as "lower-level" search result items. The search tool 154 determines whether a matching item qualifies as a top-level, as opposed to bottom level, item using criteria that will be discussed later.

[0050] For the Music category 340, no top-level search result items are displayed on the All Products search results page 300 in the FIG. 3 example. Instead, only a link to lower-level items 338 is included. This format is used when the search tool 154 finds matches within a category, but none of the matches qualify as a top-level item.

[0051] In accordance with one feature of the invention, the display order of the categories within the Top Search Results section 305 is determined by the category ranking process 150, based upon an assessment of the likely relevance of each category to the search query. Thus, in FIG. 3, the category ranking process 150 determined that, for this particular search query, the Books category 310 was likely to be of greatest relevance to the user, followed by the Videos 320, Auctions 330, and Music 340 categories. An important benefit of this feature is that it reduces the need for users to view or scroll through search results that are not of interest. The manner in which relevance is assessed will be discussed later with the help of FIG. 8.

[0052] Although only four categories are depicted (for purposes of clarity) in FIGS. 1 and 3 for products directly associated with the web site 130, the AMAZON.COM web site includes a much larger number of categories that compete for priority within the Top Search Results section 305.

[0053] Immediately below the Top Search Results section 305, the All Products search results page 300 shown in FIG. 3 displays the Additional Matches section 350. This middle section displays the results generated from application of the query to the affiliated merchant databases, that is, to the Software and Electronics databases 145, 146. These results are categorized in FIG. 3 under the Software 360 and Electronics 370 headings. Within each category up to three items associated with that category are preferably displayed to the user. For example, in the Software category the items "A Horse's Tale" 362, "Extracts from Adam's Diary" 364, and "A Visit to Heaven" 366 are the top three matches as assessed by the search tool 154. As noted above, the search tool 154 assesses whether or not an item qualifies as a top-level item. Hypertext links (368, 378) to additional matches (i.e. lower-level items) are provided for the respective categories.

[0054] The Electronics heading 370 shown in FIG. 3, like the Music category 340 above, does not display top-level results, but instead includes a link to lower-level results. Again, this presentation is used when the search tool 154 finds matches within the category, but fails to find any matches qualifying for top-level status.

[0055] The order of the categories within the Additional Matches section 350 is determined by the category ranking process 150, based upon an assessment of the likely relevance of each category to the search query. Thus, in FIG. 3, the category ranking process 150 determined that, for this particular search query, the Software category 360 was likely to be of greater relevance to the user than the Electronics category 370. Although only two categories are depicted (for purposes of clarity) for the affiliated merchant databases in FIGS. 1 and 3, the AMAZON.COM web site includes a large number of categories that compete for priority within the Additional Matches section 350.

[0056] Immediately below the Additional Matches section 350, the results page displays the Related Products section 380. This section displays the search results generated from application of the query to the unaffiliated merchant database, that is, to the Product Spider database 147. In the preferred embodiment, no top-level results are displayed for this category. Instead, the results are accessible from the All Products search results page 300 via a hypertext link labeled "Related Products" 380. The search of the Product Spider database 147 preferably does not take place simultaneously with the searches of the other databases 141-146. Rather, the Product Spider search is initiated by the user's selection of the Related Products hypertext link 380, instead of by the user's selection of the search initiation button 240.

[0057] In another embodiment, the top three (or more) Product Spider results are displayed on the All Products results page 300 together with the categories in the Top Search Results and Additional Matches sections.

[0058] The display format of the All Products search results page illustrated in FIG. 3 allows a user to very efficiently identify all of the categories of products that are

relevant to the query submitted by the user. This efficiency results in part from the limited number of items displayed to the user within each category.

[0059] FIG. 4 illustrates the general format of a Related Products search results page 400 generated in response to the selection of the Related Products hypertext link 380 in FIG. 3. The Product Spider results page 400 displays the query search result items found in the Spider database 147. The results are displayed as hypertext links in order of likelihood of relevance to the search query as assessed by the search tool 154.

[0060] As noted previously, the search tool 154 assesses the relevance of a multiple-term query to the Product Spider database 147 through inverse document frequency. That is, the weight given to a query term is inversely proportional to the frequency with which it appears in the database. For example, if a user enters the multiple-term query "Mark Twain" into the search engine query field 230, the term "Twain" is likely to appear far less than the term "Mark" in the Product Spider database 147. As such, when searching the database 147 the search tool 154 will give far greater weight to the term "Twain" when prioritizing the results for display to the user. The search tool 154 further prioritizes the results according to each query term's number of appearances, and location of appearance, within the web page. Appearances in the web page title are given the greatest priority; appearances in the first eight words of the body are given secondary priority; appearances in the subsequent thirty-two words of the body are given tertiary priority; and appearances in the remainder of the body are given lowest priority. This priority scheme, which is included with the search tool software developer's kit, is adjustable as needed.

[0061] Furthermore, as discussed below, the search tool 154 may use the product score values 170 (indicative of the likelihood that the corresponding web pages contain products available for purchase) stored in the Product Spider Database 147 to assist in the prioritization of the results generated from the Product Spider Database 147.

[0062] FIG. 4 indicates that applying the above rules to the Product Spider database 147 for the query "Mark Twain" provides five highest ranking results, the top three of which are entitled, "Mark Twain: Wild Humorist of the West" 410, "Vintage Lifestyles—A visit to Mark Twain's House" 420, and "Celebrated Jumping Frog of Calaveras County" 430. Each search result item includes a hypertext link to a web page of the unaffiliated merchant associated with the result. For example, selecting the top search result item takes the user to the unaffiliated merchant web site located at the URL: "http://207.98.171.148:80/books/twain.html" 414. Each search result item also includes a short squib 412 derived from the web page during the creation of the Product Spider database 147.

[0063] Since the number of matches found for a search query can be quite large, the results are generally partitioned so that only the top results are displayed on the first results page. In FIG. 4, for example, the first results page only displays the top five search result items. Typically the first page will display substantially more than five items. Matches with lower priority are displayed on additional Related Products results pages. These pages are accessible in sequential order via a "Next" button 440, or through a direct access link 442. Furthermore, the user may further refine the

search by accessing a refinement link 450 that allows the submission of additional search query terms.

[0064] In one embodiment, the Product Spider results page 400 includes a rating 460 for one or more of the displayed result items. In FIG. 4, a rating (three out of five stars) is associated with the fifth result item (Autobiography of Mark Twain) based upon ratings provided by users who have previously interacted with this same on-line merchant. In one embodiment, the rating information is stored as an entry in the Product Spider database 147 (not shown in FIG. 1) corresponding to the on-line merchant's URL. In another embodiment, the rating information is stored in a separate "ratings" database, indexed by canonical URL. In these embodiments, the Product Spider results page 400 further includes an option, such as a hypertext link 462, for the user to rate the merchant. The option to rate the merchant preferably exists for every result item displayed on the Product Spider results page 400, and in all lower level results pages. FIG. 4 shows such an option for only one of the five merchants merely for the sake of clarity.

[0065] Additional details regarding the presentation of the Product Spider results page are provided below following the description of FIG. 5.

[0066] B. Method for Generating a Product Spider Database

[0067] FIG. 5 illustrates the sequence of steps that are performed to construct and refresh the Product Spider database 147. In step 510, the web crawler 160 crawls a fraction X of the World Wide Web. Web crawling programs, which attempt to locate all web pages accessible on the World Wide Web by following hypertext links, are well known in the art. The size of the fraction X of the World Wide Web that is crawled in step 510 depends upon the frequency with which the Product Spider database is refreshed. The World Wide Web presently contains a sufficiently large number of web pages so as to require an extended period of time for complete crawling. As such, if the Product Spider database 147 is refreshed frequently, only a fraction X of the World Wide Web is crawled between database updates.

[0068] As shown in FIG. 5, the web pages found through step 510 are passed through a page analyzing step 520 in which the non-content based characters of the web page HTML code (e.g., the typesetting characters, the hypertext link indicators, etc.) are removed. The remaining characters correspond to the text-based content of the web page. This content is passed to the product score generator 162 which generates a numerical score between 0 and 100 indicative of the likelihood that the source web page offers a product for sale. Scores of 0 and 100 indicate the smallest and largest likelihood, respectively, that the page is offering a product for sale. Loosely speaking, the product score may be thought of as a degree of confidence written as a percentage of absolute certainty.

[0069] The analysis is conducted on a page by page basis, with each web page being assessed independently. In an alternative embodiment, a target page may be assessed by analyzing, in addition to the content of the target page itself, the contents of other web pages linked to the target page. The analysis may be limited to "neighboring" web pages (i.e., web pages directly accessible via a link on the target page), or it may extend to encompass more remotely acces-

sible web pages (i.e., web pages that are only accessible via a series of links). In these embodiments, the contributions of other web pages to the assessment of the target page may be weighted such that the influence of a remote page decreases with the number of links between the page and the target page, and/or such that only web pages of the same web site are considered.

[0070] In yet another embodiment, a web site may be analyzed as a single entity. The web site assessment may occur by combining the results of a page by page assessment of the web pages within the site, or it may occur by analyzing the web site as a whole.

[0071] The page analyzing step 520 also looks for character strings judged to be inappropriate for users of the host web site 130. For example, web sites identified as marketing salacious adult content are excluded from the Product Spider database 147.

[0072] 1. Score Generation Process

[0073] To produce a product score, the product score generator 162 first generates a set of confidence parameters designed to assess the degree to which the content-based text of a web page suggests a product is being offered for sale. One confidence parameter, "HasOfferingPrice," quantifies the presence of character strings indicative of offering prices. To create the HasOfferingPrice parameter, the product score generator 162 parses the page contents looking for character strings indicative of currency, such as "\$," "US\$" (for prices in dollars), "£" (for prices in pounds), and "dm" (for prices in Deutschmarks) followed by a string of digits. The algorithm also looks for strings indicative of an offering price, such as "price is []," "price: []," "list: []," "regularly: []," "our price is []," "price including standard shipping is []," "cost is []," "on sale now at []," "on sale now for []," and "[] for one," where in each case the square brackets signify a currency indicator followed by a string of numbers. Each time a character string denotative of an offering price is found, the HasOfferingPrice parameter is incrementally increased through a "NoisyOr" operation with a weighting factor.

[0074] The NoisyOr operation is an analog variant of the binary OR operation, where $\text{NoisyOr}(A, B) = A + B - (A \times B)$, where A and B are between 0 and 1, inclusive. The properties of the NoisyOr operation are characterized in Table I, where a variable B ($0 \leq B \leq 1$) is NoisyOr'ed against select values of a parameter A ($0 \leq A \leq 1$).

TABLE I

A	B	NoisyOr(A, B)
1	B	1
$\frac{3}{4}$	B	$\frac{3}{4} + \frac{1}{4} B$
$\frac{1}{2}$	B	$\frac{1}{2} + \frac{1}{2} B$
$\frac{1}{4}$	B	$\frac{1}{4} + \frac{3}{4} B$
0	B	B

[0075] If A=1, the output is 1 regardless of the value of B, as expected for an OR operation. If A=0, the output is B regardless of the value of B. For intermediate values of A, the output is A summed with a fraction of B. The resultant output is always equal to or larger than the larger of the two inputs, but never bigger than 1.

[0076] The weighting factor used for a particular text pattern for a particular confidence parameter within the NoisyOr operation depends upon the degree of confidence associated with the text pattern. Table II provides weighting factors, determined empirically, for the HasOfferingPrice parameter for some example text patterns. In Table II the parameters H_{old} and H_{new} refer to the HasOfferingPrice parameter before and after, respectively, the NoisyOr operation is applied.

TABLE II

Text Pattern	Weighting	Resulting NoisyOr
"your price is []"	0.9	$H_{new} = \text{NoisyOr}(H_{old}, 0.9)$
"price: [] for one"	0.6	$H_{new} = \text{NoisyOr}(H_{old}, 0.6)$
"price is [] per person"	0.4	$H_{new} = \text{NoisyOr}(H_{old}, 0.4)$
"[]/person"	0.2	$H_{new} = \text{NoisyOr}(H_{old}, 0.2)$
"[] for one"	0.1	$H_{new} = \text{NoisyOr}(H_{old}, 0.1)$

[0077] Inspection of Table II indicates that the character string "your price is []" is believed to be a very good predictor of web pages offering products for sale. The character string "[] for one," on the other hand, is believed to be a relatively weak predictor.

[0078] Other confidence parameters, analogous to HasOfferingPrice, are used to quantify a wide variety of character strings associated with product offerings, including the presence of warranty terms, sales tax information, shipping information, SKU numbers, shopping carts, and click-to-buy options. Each confidence parameter is incremented through the use of the NoisyOr operation and weighting factors in the same manner described above for the HasOfferingPrice parameter. The specific character strings and weighting factors used for the confidence parameters are disclosed in Appendices A, B and C of the incorporated-by-reference parent application.

[0079] The product score generator 162 combines the finished set of confidence parameters through a series of nested NoisyOr operations, again using empirical weighting factors, to generate a single product score for the page. The specific combinations and weighting factors used to generate the product score are disclosed in Appendix D of the incorporated-by-reference parent application.

[0080] 2. Second Analysis Stage

[0081] In practice, the vast majority of the web pages on the World Wide Web are not associated with product offerings, and as such their corresponding product scores are low. As shown in FIG. 5, these web pages are excluded from the Product Spider database 147 by a filtering step 530. The filter is simply a threshold number, preferably thirty, that the web page product score must equal or exceed to satisfy the filter. Web pages having a product score below thirty are discarded 532 as inappropriate for the Product Spider database 147. Typically about 99% of all web pages in the World Wide Web are discarded in this manner. Those pages having product scores satisfying the filter criteria are retained. The corresponding URLs are submitted back 540 to the web crawler 160 for a second crawling stage 560.

[0082] In other embodiments, such as those in which the index is also used to provide a general purpose web search engine, pages may be indexed without regard to their

respective product scores. In still other embodiments, the filter comprises multiple ranges of product score values with predetermined minimum and maximum values. For example, four separate databases may be created for web pages having product score values of 20-40, 40-60, 60-80, and 80-100, respectively. In these latter embodiments the product scores may optionally be omitted from the respective databases.

[0083] If the Product Spider database is not being constructed for the first time, but rather is being updated, then the URLs from the existing database 147 are submitted 550 to the second crawling stage 560 as well. Duplication between the previous database submissions 550 and the latest web crawl submissions 540 are detected and removed (not shown).

[0084] The second crawling stage 560 shown in FIG. 5 typically requires substantially less time than the first crawling stage 510, as the number of web pages involved is considerably smaller. The results of the second web crawling stage are passed through a second page analyzing stage 570, wherein product scores are generated anew. In a second filtering stage 580, pages failing to satisfy the filter are once again discarded 582. Those pages satisfying the second filtering stage 580 are passed in step 590 to the index tool 164 for further processing.

[0085] The second filtering stage 580 preferably uses the same criteria as the first filtering stage 530. In an alternative embodiment, the second filtering stage 580 may have either more or less discriminating criteria than the first filtering stage 530.

[0086] 3. Construction of the Product Spider Database

[0087] The pages retained after the second filtering stage 580 shown in FIG. 5 are passed to an indexing stage 590 wherein the index tool 164 creates the Product Spider database 147, fully text indexed by keyword 166. A given web page will contain multiple index keywords distributed throughout its text. The index tool 164 converts the information from a form organized by URL into a form organized by keyword. Schematically, the index tool 164 reorganizes the set of multiple pages (Page_m , where $m=1$ to M) containing multiple Keywords (Word_n , where $n=1$ to N) such that $\text{Page}_1(\sum_n \text{Word}_n)$, $\text{Page}_2(\sum_n \text{Word}_n)$, \dots , $\text{Page}_M(\sum_n \text{Word}_n)$ is converted into $\text{Word}_1(\sum_m \text{Page}_m)$, $\text{Word}_2(\sum_m \text{Page}_m)$, \dots , $\text{Word}_N(\sum_m \text{Page}_m)$.

[0088] As shown in FIG. 1, the database 147 includes, for each keyword 166, one or more web page addresses 167 with corresponding titles 168, squibs 169, and product scores 170. All of the product scores will necessarily equal or exceed thirty in the preferred embodiment due to the second filtering stage 580.

[0089] The web page addresses 167 stored in the Product Spider database 147 are preferably "canonicalized" URLs. URLs often include one or more strings of characters appended to the addressing information that specify, for example, a particular user ID, session ID, or transaction ID. These characters are not needed for accessing the web page, and are thus preferably discarded, resulting in a "canonical" URL for inclusion in the Product Spider database 147. Techniques for canonicalizing URLs are well known in the art.

[0090] The title 168 entry of the database 147 is preferably duplicated directly from the title used for the web page, as identified by the appropriate HTML tags. If a web page has an inappropriate title, or is missing a title, a new title is inserted into the database 147 as needed on a case by case basis.

[0091] The squib 169 entry of the database 148 is generated automatically by the index tool 164. The squib corresponds to the initial series of words on a web page, up to a preset number of characters set at about two-hundred. In another embodiment, the squib displays relevant text extracted from the web page corresponding to the products offered for sale on the web page.

[0092] The process illustrated in FIG. 5 may be used to update the Product Spider database 147 as often as desired. In a preferred embodiment, the Product Spider database 147 is updated every week, more preferably the database is updated every three or four days, and even more preferably it is updated every day.

[0093] As indicated above, the Product Spider database 147 may alternatively be constructed without storing the product scores for each page. In one embodiment, for example, the database comprises only pages having a product score satisfying predetermined criteria, for example, requiring the product score to equal or exceed thirty (as in the filtering steps 530, 580 of FIG. 5). In another alternative embodiment, the database comprises multiple indexed tables created without storing the product scores, wherein each table is constructed from web pages having a product score satisfying unique criteria, for example, four separate indexed tables containing pages having product scores from 20-40, 40-60, 60-80 and 80-100, respectively.

[0094] In another embodiment, the Product Spider database 147 consists of multiple indexed tables, wherein each table is constructed from web pages that are distinguishable on the basis of some aspect of product offerings (ascertained from parsing the web pages) unrelated to product scores. In one embodiment, for example, the database 147 consists of separate tables for different categories of goods (e.g., books, music, videos, electronics, software, and toys). In another embodiment, a separate table is used for products unsuitable for children. In still another embodiment, different tables are constructed for web sites written in different languages (English, Japanese, German, etc.). In yet another embodiment, different tables are constructed for on-line and off-line product offerings. Under these embodiments, the page analyzer steps 520, 570 include searching for character strings judged to be associated with the various predefined categories.

[0095] By constructing the Product Spider database 147 out of different tables having distinguishing characteristics, or retaining the equivalent information within one big table, the user is capable of conducting a more refined search within the Product Spider database 147. In one embodiment, for example, the Related Products hypertext link 380 is replaced by a pulldown menu comprising different categories corresponding to the distinctions retained within the Product Spider database 147 (e.g., books, music, video, and toys categories, on-line versus off-line offerings, goods versus services, etc.).

[0096] The Related Products search results page 400 (FIG. 4) displays search results using TFIDF prioritization

as applied to the entire Product Spider database 147. That is, the results consist of a single list drawn from all of the web pages satisfying the filtering steps 530, 580 of FIG. 5 (e.g., all web pages having product scores at or above 30). In another embodiment, the search results are presented as a number of lists, with each list having an independent TFIDF prioritization. In this embodiment, each of the multiple lists consists of pages satisfying different product score criteria. In one embodiment, these multiple lists are displayed separately. In an alternative embodiment, the lists are concatenated into one long list. This latter embodiment is illustrated in Table III, where an All Products search of the host web site 130 generates two Product Spider results lists, an "A" list for web pages having product scores at or above eighty, and a "B" list for web pages having product scores below eighty but at or above thirty. The A List is presented first (on multiple pages if it is long), and the B List is concatenated onto the end of the A List. This bifurcation attempts to provide the user first with those pages most likely to be offering a product for sale. The B list may optionally be generated only if the number of items in the A list falls below a threshold number, such as five, or if the user requests to view (e.g. via selection of a hypertext link) the B list.

TABLE III

Lists	Criteria
A List	Product Score ≥ 80
B List	$80 > \text{Product Score} \geq 30$

[0097] The Product Spider search feature has been discussed above in the context of a product-oriented search engine of an on-line merchant. The feature may be implemented in other contexts as well. It may be implemented, for example, as part of a "general purpose" web search engine as a user-selected option (e.g. through a pulldown menu or through selection of a "product offering" button).

[0098] C. Method for Ranking Categories in an All Products Search

[0099] As noted above, users of the AMAZON.COM web site 130 may conduct an All Products search that will generate results for items directly offered for sale by the AMAZON.COM web site (organized into multiple categories), items offered for sale by third parties (Auction and zShop users) using the Amazon web site as a forum, items offered for sale by other on-line merchants affiliated with AMAZON.COM (organized into multiple categories), and items offered for sale by on-line merchants unaffiliated with AMAZON.COM (those within the Product Spider database 147). With such a large number of categories involved, it is advantageous that the results of such a cross-category search be displayed efficiently. In particular, it is desirable that the search results of most relevance to the user be displayed so that the user does not need to wade through a long list of irrelevant search results or click through a long series of hypertext links to find the results of greatest interest.

[0100] FIG. 6 illustrates the sequence of steps that are performed to construct an All Products search results page such as depicted in FIG. 3. In a first step 610, the user is prompted to enter a search query to all of the products

databases 141-147. One approach is illustrated in FIG. 2, where the user may select the All Products option 260 from the pulldown window 250.

[0101] In a second step 620 shown in FIG. 6, the user submits a search query to be applied to all of the product databases 141-147. One approach is illustrated in FIG. 2, where the user may enter a query into the query field 230 and select the search initiation button 240.

[0102] In a third step 630 shown in FIG. 6, the query is applied to all of the categories comprising the All Products search. As illustrated in FIG. 1, the query is submitted by the query server 140 via the search tool 154 to a separate database 141-147 associated with each category. Each product database 141-147 is indexed by keyword to facilitate searching by the search tool 154.

[0103] In a fourth step 640 shown in FIG. 6, the query results are returned from each of the product databases 141-147, via the search tool 154, to the query server 140. The search tool prioritizes the results within each category according to a determination of relevance based upon the query terms. The method used by the search tool 154 to prioritize the search result items within a category varies depending upon the nature of the category searched. The prioritization methods used are discussed at length below with the help of FIG. 7.

[0104] In a fifth step 650, a relevance ranking is generated for each competing category based on an assessment of the relevance of the search query to that category. The method used to generate the relevance ranking is discussed in more detail below with the help of FIG. 8.

[0105] In a sixth step 660, the categories are arranged in a display order determined by the results of the category relevance ranking step 650. The primary purpose of this step is to display the categories (and associated search results) deemed to be the most closely related to the search query near the top of the search results page. In a final step 670, a search results page having the appropriate arrangement is generated for display to the user. An example All Products search results page 300 is illustrated in FIG. 3. The approach discussed above displays to the user the top-level search results deemed to be of greatest interest to the user in a manner that is efficient (long lists are avoided), comprehensive (all pertinent categories are included, and links to further results are provided), and clear (the organization and prioritization helps the user quickly comprehend the results).

[0106] 1. Prioritizing Search Result Items Within Each Category

[0107] The method used to prioritize the search result items within the Product Spider database 147 has already been discussed. Briefly, a database entry is given a higher priority depending upon the number of times a search term appears in the page. Appearances in the web page title, and in text near the beginning of the page, are given higher priority than later appearances. In multiple-term queries, the significance of each term is weighted in a manner inversely proportional to how frequently the term appears in the Product Spider database 147.

[0108] Search result items within the Auctions database 144 are prioritized based upon the ending time of the auction, with a shorter closing time receiving higher priority

than a later closing time. The top-level Auctions result items displayed on the All Products search results page (see FIG. 3) correspond to the matching Auctions result items (preferably up to a maximum of three) having the most imminent ending times. Selecting the hypertext link to the lower-level matches 338 provides a list of all of the matching Auctions result items. These lower-level results may be sorted by the auction ending times, by the TFIDF relevancy of the search query, by the starting time of the auction, by the present number of bids, or by the highest present bid. The method of display of these lower-level results is preferably an option for the user (e.g. via a pulldown menu).

[0109] For databases associated with the AMAZON.COM zShops (fixed-price offerings by third parties—not shown in FIG. 1), the top-level search result items are prioritized in the same manner used for the Product Spider database 147, discussed above. Thus the zShops results displayed on an All Products search results page correspond to the zShops matching result items (preferably up to a maximum of three) having the highest TFIDF relevancy. Selecting an associated hypertext link to lower-level matches provides a list of all of the matching zShops result items. These lower-level results may be sorted by TFIDF relevancy, by the starting or ending date of the zShop, or by the product price. The method of display of these lower-level results is preferably an option for the user (e.g. via a pulldown menu).

[0110] For items within the categories of goods sold directly by the host web site (i.e. goods from the Books, Music, and Videos databases 141, 142, 143) or sold by affiliated merchants (i.e. goods from the Software and Electronics databases 145, 146) the results are prioritized using a more sophisticated approach than those discussed above. For these categories, the assessed relevance of a search result item is based upon the frequency with which the item has been selected in the past during similar queries. The manner in which this is accomplished is now discussed with reference to FIG. 7.

[0111] FIG. 7 illustrates the structure of the Books database 141 of FIG. 1. The database consists of two tables, a Books Full Text Index 710 and a Books Popularity Score Table 750. The Books Full Text Index 710 contains information, indexed by keyword, for every item in the Books catalog of the web site 130. The Books Popularity Score Table 750 contains information about the subset of books from the Books catalog that users of the web site 130 have recently “selected” during on-line searches.

[0112] The Books Full Text Index 710 is indexed by keyword 712 to facilitate searching by the search tool 154. The comprehensive indexing is created in a manner analogous to that discussed above for the Product Spider database. Briefly, the index tool 164 converts the information from a form organized by item into a form organized by keyword. Thus, for each keyword 712, the Books Full Text Index 710 contains one or more item identifiers 714 each of which uniquely identify a book within the on-line catalogue of the host web site 130.

[0113] The Books Full Text Index 710, for example, associates the keyword “Twain” 720 with eight distinct item identification numbers, each corresponding to a single book. Inspection of FIG. 7 reveals that the term “Twain” is associated with a book corresponding to item identification code 1311302165. This association between a keyword and

a book may come from the word appearing in the book's title, in the author's name, or in ancillary text, such as descriptions and third party reviews of the book.

[0114] The Books Popularity Score Table 750 is also indexed by keyword 752. For each keyword 752, the table 750 contains one or more item identifiers 754 analogous to those of the Full Text Index 710. The table also includes, for each keyword-item pair, a "popularity score" 756, the meaning of which is discussed below.

[0115] The entries in the Books Popularity Score Table 750 are generated through the actions of users conducting on-line product searches on the web site 130. When a user conducts a search within the host web site 130, the user's search query is stored in the query log 136 shown in FIG. 1. The hypertext links selected by the user following the search are also stored, as are the times at which the selections are made. Through parsing of the query log 136, the user's actions may be followed in great detail. A query log parsing processor (not shown in FIG. 1) extracts the relevant information and generates the popularity scores 756 stored in the Books Popularity Score Table 750. The manner in which this is done, together with more details about using popularity scores to facilitate query searches, is described in the U.S. application Ser. No. 09/041,081, filed Mar. 10, 1998, entitled "Identifying the Items Most Relevant to a Current Query Based on Items Selected in Connection with Similar Queries," which is hereby incorporated by reference.

[0116] The popularity scores 756 of the Books Popularity Scores Table 710 reflect the frequency with which users have selected the corresponding item 754 from query results produced from searches containing the corresponding keyword 752 as a query term. For example, FIG. 7 indicates that the item associated with code 2722601080 has been selected by a user one time following a search including the query term "Mark"760. The item identified with code 4603283881, by comparison, has been selected twenty-two times following searches including the query term "Mark"760. This latter book has also been selected forty-one times following searches including the query term "Twain"770.

[0117] Different actions by a user may be used to qualify as a "selection" for purposes of determining the popularity score 756. Actions may include, for example, displaying additional information about the item, spending certain amounts of time viewing information about the item, accessing hypertext links within the information about the item, adding the item to a shopping basket, and purchasing the item. All of these actions may be assessed from the query log 136.

[0118] Different weightings may be associated with different user activities. In one embodiment, for example, clicking on an item increments the item's popularity score by one while placing the same item in an on-line "shopping cart" increments its popularity score by fifty.

[0119] Preferably the popularity scores 756 are determined by the recent actions of users over a predetermined amount of time, such as a week, ensuring that the scores represent current user preferences. Preferably the Books Popularity Score Table 750 is constructed by merging the results of a number of intermediate tables corresponding to user actions over adjacent periods of time. The query log 136

is parsed once per day to generate a daily intermediate table containing keyword-item pairings and corresponding popularity ratings for that day. Each day, a new full table 750 is constructed by merging the new intermediate table with the most recent N intermediate tables, where N is a predetermined number. The parameter N is selected to equal thirteen for all categories. This creates a full table 750 representing results over a "sliding window" in time fourteen days in duration. In another embodiment, the number N is selected to be larger for categories that experience low user traffic (e.g. Classical Music) than for categories that experience high user traffic (e.g. Books).

[0120] The popularity scores of the multiple intermediate result tables are weighted equally during the merging into the full table 750. In an alternative embodiment, the popularity scores of the multiple intermediate tables are assigned different weightings for merging, with the weightings depending on the times at which the intermediate tables were created. In one such embodiment, the weightings used for merger decrease with increasing age of the intermediate table.

[0121] If a book falls out of fashion, and thus is not selected within the time period stored in the present version of the Books database 141, the book will fail to appear in the associated Books Popularity Score Table 710. If a book has been selected within the relevant time period, it will contain one entry in the table 710 for every query term utilized by users prior to selecting the book over that time period. During a later period the same book may have completely different keyword 712 entries if the users selecting the book utilized different search query terms to find it.

[0122] All Products search results from the Books database 141 are prioritized, for purposes of display on the All Products search results page 300, based on the popularity scores 756 of the Books Popularity Score Table 750.

[0123] For single-term queries, the search tool 154 prioritizes the search result items based upon each item's popularity score. Referring to FIG. 7, for example, a search of the Books database 141 for the query term "Mark"760 would prioritize the three items as 4603283881, 9040356769, and 2722601080, as determined by the three popularity scores 22, 7, and 1, respectively.

[0124] For multiple-term queries, such as "Mark Twain," the search tool 154 only returns items having entries in the Books Popularity Score Table 750 under both query terms. In FIG. 7, for example, the query "Mark Twain" would trigger a match for item 4603283881, which is present for both query terms, but not for the other items displayed. When multiple items match all of the terms of a multiple-term query, the popularity scores of each term for that item are combined in some manner to create a query phrase popularity score for that item. In one embodiment, the query phrase popularity score is the sum of the popularity scores 756 of the component terms. In other embodiments, discussed later, a more complicated combination of the scores is used.

[0125] This prioritization scheme is used to determine the top-level matches that are displayed on the All Products search results page 300. The top-level matches correspond to those items, up to a maximum of three, having the highest popularity scores for the submitted query. For example, in

FIG. 3 the three top-level matches **312**, **314**, **316** under the Books category **310** represent the three highest popularity scores for the search phrase "Mark Twain." Furthermore, the top-level matches are ordered on the All Products search results page **300** based on popularity score. Referring to **FIG. 3**, the items labeled "Letters from the Earth"**312**, "Following the Equator . . ." **314**, and "Joan of Arc"**316**, had the first, second, and third highest popularity scores, respectively, in the Books database **141** for the submitted search query.

[0126] For the Books category **310** of an All Products search, the lower-level matches are accessible from the All Products search results page **300** via a hypertext link **318**. This link **318** generates a lower-level Books results page that displays both the top-level search result items and lower-level search result items.

[0127] In one embodiment, the lower-level search result items from the Books Popularity Score Table **750** that matched the submitted query are displayed most prominently, followed by search result items found only in the Full Text Index **710**. In an alternative embodiment, the lower-level search result items are displayed according to preset categories unrelated to popularity score **756**. In one embodiment, for example, the lower-level Books results page may display three separate alphabetized lists: one for books that are immediately available, one for books that must be special ordered, and a third for books that are currently out of print. Preferably the user is provided with the ability to search the Books Full Text Index **710** based on other criteria as well, such as author, title, and ISBN (International Standard Book Number).

[0128] The Music, Videos, Software, and Electronics databases **142-146** are structured in the same manner as the Books database **141** shown in **FIG. 7**. Each category database consists of a Full Text Index, containing comprehensive information about the products within the category, and a Popularity Score Table containing information about recent search and selection activities by users within the category. The Auctions database **144**, on the other hand, includes an Auctions Full Text Index but lacks an analog to the Popularity Score Table.

[0129] Upon initiation of an All Products search query, the search tool **154** returns a prioritized list of search result items, for each category, using the approach discussed above for the Books category. The top matches from this prioritized list (up to a maximum of three) become the "top-level" matches, for each category, for display in the All Products search results page **300**. For these categories **320**, **330**, **340**, **360**, **370**, as for the Books category **310**, lower-level search result items are accessible from the All Products search results page **300** via a hypertext link **328**, **338**, **348**, **368**, **378**.

[0130] There are three special circumstances in which prioritization of one or more of the Books, Music, Video, Software, and Electronics categories is not based on popularity scores. First, when a category is newly introduced it takes some time (e.g. two weeks) for the corresponding Popularity Score Table to accumulate sufficient user selections to result in useful popularity scores. During this "introductory period" the top-level All Products search result items are prioritized using the TFIDF relevancy approach discussed earlier.

[0131] The second special circumstance arises when a new product line is introduced within a category. In this situation

popularity scores may be plentiful for the category as a whole, but scores for the newly released product line will necessarily be lacking. In order to assist users in finding the new product line during this "transition period," the top-level All Products search result items for that particular category are prioritized using the TFIDF relevancy approach discussed earlier.

[0132] The third special circumstance arises when a search query is so unusual that the search tool **154** fails to generate a single match within any of the Popularity Score Tables (or in the Auctions database **144**). This circumstance is discussed at length after the following section with the help of **FIG. 9**.

[0133] For some All Products search queries, the search tool **154** may find matches within the Books Popularity Score Table **750**, but not within the analogous Music Popularity Score Table. In this case, there are no top-level search results available to display on the All Products search results page **300** for the Music category **340**. Indeed, this is what is displayed in **FIG. 3**. If the search tool **154** finds at least one lower-level result item (i.e. a match in the Music Full Text Index), a hypertext link **348** to the lower-level results is provided on the All Products search results page **300**. Inspection of **FIG. 3** reveals that this is the case for the Music and Electronics categories **340**, **370** in response to the query "Mark Twain." If, on the other hand, no matches are found in either the category's Popularity Score Table or Full Text Index, then the category is omitted entirely from the All Products results page **300**.

[0134] 2. Ranking Categories Based on Relevance

[0135] Once the search tool **154** has generated search results from each of the categories, the categories themselves compete for priority for display purposes. These competitions between categories, like the ranking of items within each category, are based upon an assessment of the relevance of the search query to each competitor.

[0136] **FIG. 8** illustrates the category ranking process **150** used to generate a category relevancy ranking for each competing category in an All Products search. The categories involved in an All Products search do not all necessarily compete with one another. Rather, the categories may be divided up into a number of "sets." Within each set, the member categories compete for priority for display purposes. Categories from different sets do not compete. Different sets might themselves compete for priority, or their arrangements may be predetermined. In different embodiments, the sets may be grouped into "sets of competing sets," and so on, as needed.

[0137] The categories of the host web site **130** are divided, for purposes of an All Products search, into three sets of categories. These sets are most easily seen through inspection of the All Products search results page **300** in **FIG. 3**. One set consists of categories that compete for priority within the Top Search Results **305** section of the results page **300**. Another set consists of categories that compete for priority within the Additional Matches **350** section of the results page **300**. A third set consists of a single category, the Product Spider results that are accessible through the Related Products **380** hypertext link.

[0138] Referring to **FIG. 8**, in a first step **810** the query server **140** identifies a first set of competing categories. The

query server 140 may identify, for example, the set of categories competing for display space in the Additional Matches 350 section of the All Products search results page 300. These categories are exemplified in FIG. 1 by the Software and Electronics databases 145, 146.

[0139] In a second step 820, the query server 140 examines the search results for a first category within the first set. The query server 140 may examine, for example, the top-level search result items for the Software category. The first column of Table IV provides an example of All Products search result items determined from the Software database 144 for the search query "Mark Twain." The search tool 154 determined that the three best top-level Software Category result items are "A Horse's Tail," "Extracts from Adam's Diary," and "A Visit to Heaven" (these results are also displayed in FIG. 3). The number in parenthesis adjacent to each item represents the popularity score for that item (see FIG. 7 and associated discussion) for the search query "Mark Twain."

TABLE IV

Software	Flowers & Gifts	Packaged Travel
A Horse's Tail (59)	Mark Twain Riverboat (57)	Autumn in the Ozarks (61)
Extracts from Adam's Diary (20)	On the Trail of Mark Twain (13)	Bermuda (4)
A Visit to Heaven (11)		Europe—Atlantic Crossing (1)

[0140] In a third step 830, the query server 140 determines a category "popularity" score indicative of the significance of the query term to the category. The category popularity scores are generated from some aspect (e.g., the popularity scores) of the constituent search result items in each category. In one embodiment, the category popularity score is determined by summing the constituent top-level result item popularity scores. Applying this approach to the Software results shown in Table IV leads to a category popularity score of 90 (=59+20+1).

[0141] In a fourth step 840, the query server 140 repeats the above examination of the All Products search results for another category within the first set. For example, the second and third columns of Table IV show search results for Flowers & Gifts and Packaged Travel categories (not shown in FIGS. 1 and 3) that compete with the Software category for priority within the Additional Matches 350 section of the All Products search results page 300. Determining category popularity scores using the approach discussed above results in scores of 70 (=57+13) for Flowers & Gifts and 66 (=61+4+1) for Packaged Travel.

[0142] After category popularity scores have been determined for all members of the set, a category ranking is created in a fifth step 850 based upon the relative values of those category popularity scores. The rankings are determined through a comparison of each category popularity result. For example, using the category popularity results determined above, the Software (score=90), Flowers & Gifts (score=70), and Packaged Travel (score=66) categories would be ranked first, second, and third, respectively. The categories would be arranged (boxes 660 and 670 in FIG. 6) appropriately based on this ranking. That is, the Software category results would be displayed in the most prominent

manner, the Flowers & Gifts category results would be displayed in the next most prominent manner, and the Packaged Travel category results would be displayed in the least prominent manner of the three.

[0143] In another embodiment, a set of weighting factors is applied to the set of category popularity scores. Such weighting factors may be used to help or hinder particular categories as desired. For example, if it was decided that during the holiday season the Flowers & Gifts category should be provided a competitive advantage, that category may be given a weighting factor of two, with each of the remaining categories having a weighting factor of one. With such a weighting set, the Software (score=1×90=90), Flowers & Gifts (score=2×70=140), and Packaged Travel (score=1×66=66) categories would now be ranked second, first, and third, respectively. These weighting factors may be influenced by the profile of the user who submitted the search query. Furthermore, the popularity scores may be influenced by the profile of the user who submits the search query. For example, the complete history of selections made by the user within the host web site 130 may be retained in a database (not shown in FIG. 1). This information may be used to adjust the weightings to further individualize the presentation. If the user has made 90% of her prior purchases on the host web site 130 from the Videos database 143, for example, the Videos category popularity scores may be given greater weight to reflect this individualized history.

[0144] In another embodiment, the category popularity score is determined by taking the mean value of the constituent top-level result item popularity scores. Applying this approach to the results shown in Table IV leads to category scores of 30 (=90/3) for Software, 35 (=70/2) for Flowers & Gifts, and 22 (=66/3) for Packaged Travel. Thus, under this approach, the Flowers & Gifts category results would be displayed in the most prominent position.

[0145] In another embodiment, the category popularity score is determined by taking the highest value of the constituent top-level result item popularity scores. Applying this approach to the results shown in Table IV leads to category scores of 59 for Software, 57 for Flowers & Gifts, and 61 (=66/3) for Packaged Travel. Thus, under this approach, the Packaged Travel category results would be displayed in the most prominent position.

[0146] In still another embodiment, the category popularity score is determined by combining the popularity scores of all matching items found in the category's Popularity Score Table, rather than just the matching items having the three highest popularity scores. Other manners of combining top-level result item popularity scores into category popularity scores will be apparent to those skilled in the art.

[0147] Inspection of FIG. 3 reveals that the Software category 360 "won" the competition against the Electronics category 370. This is unsurprising considering that the Electronics category 370 does not include any top-level search result items. This indicates that there were no Electronics search result items with popularity scores for the query "Mark Twain," and the Electronics category popularity score using the above embodiments would equal zero.

[0148] Referring to FIG. 8, in a sixth step 860 the query server 140 identifies another set of competing categories and repeats steps two through five 820-850. For example,

another set of the host web site 130 consists of those categories competing for display space within the Top Search Results 305 section of the All Products search results page 300. These categories are exemplified in FIG. 1 by the Books, Music, Videos, and Auctions databases 141, 142, 143, 144.

[0149] Categories in this set are handled in much the same manner as was discussed above for the categories of the previous set. For the Books, Music, and Videos categories 310, 320, 340, for example, the category popularity scores are determined from the constituent top-level item popularity scores using one of several possible approaches, as discussed above.

[0150] A complication arises, however, since the Auctions category uses a completely different approach than the other categories in determining the top-level search result items. In particular, the Auctions database 144 does not include popularity scores. Rather, as is discussed above, the highest priority top-level matching results are determined based on the amount of time remaining for each matching item's auction. The category popularity score for the Auctions category is therefore determined in a manner distinct from the other categories.

[0151] In one embodiment, the Auctions category popularity score is determined by summing up the number of matching items found by the search tool 154 for the submitted search query within the Auctions database 144. In another embodiment, the Auctions category popularity score is determined by summing up the number of matching auctions with less than a predetermined amount of time remaining. In yet another embodiment, the Auctions category popularity score is determined by a weighted summation of the number of matching auctions, with the weighting factor for a particular auction determined by the amount of time remaining for that auction. Preferably this weighting factor is inversely proportional to the time remaining for the auction.

[0152] In one embodiment, the category popularity scores for all of the categories in a competing set (including the non-Auction categories) are based upon the number of items matching the submitted search query. In another embodiment, the category popularity scores for all of the categories in a set are based upon the fraction of items in the category that match the submitted search query (i.e., the number of items in the category that match the search query divided by the total number of items in the category).

[0153] The use of category popularity score weighting factors, discussed above, is preferably used to "normalize" the popularity scores between the Auctions and the other categories. In one embodiment, the Books, Music, Auction, and Videos category 310, 320, 330, 340 popularity scores are weighted equally. In another embodiment, the Auctions category popularity score is given a weighting three times as large as the scores of the remaining categories. In still another embodiment, the Auctions category popularity score is given a weighting one-third as large as the scores of the remaining categories.

[0154] Inspection of FIG. 3 reveals that the Books category 310 is the highest on the All Products search results page 300. This indicates that the Books category 310 "won" the competition against the Videos, Auctions, and Music categories 320, 330, 340.

[0155] Although the All Products search results page 300 depicted in FIG. 3 associates rank with vertical location within a section of a web page, there are other ways in which the results of a category may be given greater priority. For example, the web page may indicate priority through the use of a different font size, or a different color, through location within a web page (as in FIG. 3), through location on separate web pages, through "framing," or by display of category relevance scores or ranking (optionally expressed as a percentage, as a number of stars, etc.). Numerous possibilities would be apparent to one skilled in the art.

[0156] The final set of categories utilized by the host web site 130 consists of a single category, the Product Spider results. These results are not displayed on the All Products search results page 300, but rather are accessible on the page 300 through the Related Products 380 hypertext link. Since this set consists of only one category, there is no competition between categories, and the relevance ranking process of FIG. 8 is not followed.

[0157] The above discussion describes the category ranking process in the context of searching for product offerings. The process is also applicable to other contexts. For example, a user searching for journal articles may be provided with a top-level search results page with a limited number of items displayed within each of multiple categories. A user searching for court opinions may be provided with results divided into state appellate opinions, federal appellate opinions, etc. A user searching for discussion groups may be provided with a search results page with the items arranged by the age of participants, subject matter of the discussion, etc. A user searching for recipes may be provided with a search results page with the items arranged by food type. A user searching for movie reviews may be provided with a search results page with the items arranged by the nature of the reviewer (syndicated newspaper columnist, amateur reviewer, etc.). Numerous possibilities would be apparent to one skilled in the art.

[0158] 3. Handling Uncommon Search Queries

[0159] For uncommon search queries, the search tool 154 may fail to find matches within the Auctions database 144 or within one or more of the Popularity Score Tables of the other databases 141, 142, 143, 145, 146.

[0160] If results are found within the Auctions database 144, the top-level results will preferably be displayed on the All Products search results page 300 regardless of whether any of the other categories have top-level results (i.e., have matching results in the Popularity Score Tables of their product databases). Similarly, as long as at least one of the non-Auction categories finds matches within the category's Popularity Score Table, those top-level results will preferably be displayed on the All Products search results page 300 regardless of whether any other category found top-level results. If no categories find top-level results, the query server 140 does not generate an All Products "no results" page. Instead, the query server 140 undertakes additional steps in an attempt to generate search results from the query for display on the All Products search results page 300. The process used by the query server 140 in this endeavor is illustrated in FIG. 9.

[0161] As represented in FIG. 9, an All Products search begins with a search of the Auctions database 141 (box 910)

and a search of the Popularity Score Tables of each of the Books, Music, Videos, Software, and Electronics categories (box 920). This is also represented in FIG. 6 by box 630. The query server 140 determines whether any top-level results are returned from any categories (box 915). If at least one result is returned, the query server 140 jumps to box 650 in FIG. 6 (box 920), and the steps discussed earlier for ranking categories are followed.

[0162] If no top-level results are returned as determined by box 915, the search tool 154 conducts a search of the Full Text Indexes of each of the Books, Music, Videos, Software, and Electronics categories (box 930). The query server 140 determines whether any results are returned from any categories (box 925). If at least one result is returned, the query server 140 process jumps to box 650 in FIG. 6 (box 920), eventually resulting in the generation of an All Products search results page 300. In this case, the top-level result items from the Full Text Indexes (preferably up to a maximum of three) are determined by TFIDF relevancy score. In one embodiment the category popularity score for each category returning results is determined from the number of matching items found for that category. In another embodiment the category popularity score is determined by the fraction of items in the category that match the submitted search query (i.e., the number of items in the category that match the search query divided by the total number of items in the category).

[0163] If no results are returned as determined by box 935, the spell checker 152 attempts to find misspellings within the submitted search query (box 940). If the spell checker 152 fails to identify any misspelled query terms (box 945), a search "no results page" is generated (box 970), notifying the user of the lack of results for the submitted search query. If the spell checker 152 successfully identifies a potentially misspelled query term (box 945), the spell checker 152 creates a new query phrase by substituting, for the potentially misspelled word, a word found in a dictionary or lookup table. The search tool 154 then repeats the process from boxes 910 through 935, as needed, using the new query phrase (box 950). If the new query phrase generates results as assessed in boxes 915 or 935, the query server 140 jumps to box 650 in FIG. 6 (box 920) and an All Products search results page 300 is generated using the substituted query. The results page 300 notifies the user that the submitted query failed to produce an exact match, and displays the modified query.

[0164] If no results are returned for the modified query as determined by box 955, the query server 140 divides the query phrase into multiple single term queries. For example, the submitted four-term query, "Twain Sawyer Becky Thatcher," which will normally only generate results if all four terms are associated with a single item, is divided into four separate one-term queries, "Twain," "Sawyer," "Becky" and "Thatcher." The query processor 140 then repeats the process from boxes 910 through 935, as needed, one time for each one-term query (box 960). Matching result items of the one-term queries compete with one another (e.g., based on popularity score in the Books Popularity Score Table 710) in the same manner as the results within a multiple-term query. In this situation, however, a priority "booster" is added to the popularity scores of result items that match two or more of the search terms. The size of the booster is given by $1,000,000 \times (N-1)$, where N is the number of terms matched.

Table V illustrates an example of the use of boosters for the four-term search query given above.

TABLE V

Item:	Matching Terms:	Item-Term Popularity Scores	Multiple Term Booster	Item-Query Popularity Score
A	Twain	3566	0	3566
B	Twain	1140	1,000,000	1,001,332
	Sawyer	192		
C	Twain	20	2,000,000	2,000,040
	Becky	8		
	Thatcher	12		

[0165] Items A, B, and C in Table V match one, two, and three of the query search terms, respectively. Without the use of boosters, Item A would be prioritized first based upon the large popularity score associated with the "Twain"-Item A pairing. Each query term, however, is considered to add discriminating value for the purpose of locating items wanted by the user. Thus the booster is used to elevate those items containing more discriminating information (i.e. more query terms). Item B, which matches two query terms, is given a 1,000,000 booster. Item C, which matches three query terms, is given a 2,000,000 booster. In this way the items that are the closest matches to the full submitted query (Items B and C in this case) are given top priority. Thus the three items of Table V would be displayed on the All Products search results page 300, under the appropriate category, in the order C, B, A, as determined by their respective Item-Query popularity scores.

[0166] In order to maintain proper normalization, the same multi-term booster values are used for searches of the Auctions database 141 and the Full Text Indexes of each category.

[0167] If any of the one-term queries generate matching results (box 965), the query server 140 jumps to box 650 and an All Products search results page 300 is generated using the multiple one-term queries. The results page 300 notifies the user that the results are merely close matches to the submitted query.

[0168] If no results are returned for the multiple one-term queries (box 965), a search "no results page" is generated (box 970), notifying the user of the lack of results for the submitted search query.

[0169] Although this invention has been described in terms of certain preferred embodiments, other embodiments that are apparent to those of ordinary skill in the art are also within the scope of this invention. Accordingly, the scope of the present invention is intended to be defined only by reference to the appended claims.

[0170] In the claims, which follow, reference characters used to denote process steps are provided for convenience of description only, and not to imply a particular order for performing the steps.

What is claimed is:

1. A method of assisting users in locating items that are arranged within a database system by category, the method comprising:

monitoring user actions performed with respect to specific items in the database system to generate item usage data;

calculating item scores for specific items in the database system, wherein an item score for an item is dependent upon the item usage data associated with that item;

receiving a search query specified by a user;

identifying, within each of multiple categories of the database system, items that are responsive to the search query ("responsive items");

for each of the multiple categories, calculating a respective category score based at least in-part on the item scores of responsive items in that category; and

determining an order in which to present the multiple categories to the user such that the order is dependent upon the category scores.

2. The method of claim 1, further comprising outputting the multiple categories in said order, together with corresponding responsive items within each of the multiple categories, for presentation to the user.

3. The method of claim 2, wherein the step of outputting the multiple categories in said order comprises generating a search results page in which each of the multiple categories is listed together with a selected subset of the responsive items falling within that category.

4. A search results page generated according to the method of claim 3 and represented within a computer.

5. The method of claim 1, wherein the category scores are calculated based further upon a profile of the user such that the order in which the categories are presented is personalized for the user.

6. The method of claim 1, wherein the category score for a given category is calculated based on the item scores of a most highly ranked subset of the responsive items within the given category.

7. The method of claim 1, wherein the category score for a given category is based further upon a number of responsive items falling within the given category relative to a total number of items falling within the given category.

8. The method of claim 1, wherein the item usage data reflects item viewing events.

9. The method of claim 8, wherein the item usage data further reflects the following types of user actions: selections of items to add to a shopping cart, and item purchases.

10. The method of claim 1, wherein the item scores indicate current popularity levels of items among a population of users of the database system.

11. The method of claim 1, wherein each category score represents a predicted level of significance of the respective category to the search query.

12. A server system that embodies the method of claim 1.

13. A computer-implemented method for processing a search query specified by a user, the method comprising:

identifying, within each of a plurality of categories, one or more items that are responsive to the search query ("responsive items");

calculating a respective category score for each of the plurality of categories, wherein the category score for at least one of the categories is dependent upon a number of responsive items within that category relative to a total number of items in that category; and

presenting the plurality of categories to the user, together with associated responsive items, in an order which depends upon the category scores.

14. The method of claim 13, wherein the category score for at least one of the categories is dependent upon item scores of responsive items falling therein.

15. A web page generated according to the method of claim 13 represented within a computer, said web page displaying responsive items arranged by category in said order.

16. A server system that processes search queries according to the method of claim 13.

17. A computer readable medium which stores a computer program that embodies the method of claim 13.

18. A computer-implemented method for processing a search query specified by a user so as to personalize a presentation of search results, the method comprising:

identifying, within each of a plurality of categories, one or more items that are responsive to the search query ("responsive items");

calculating a respective category score for each of the plurality of categories, said category score reflecting a predicted level of affinity the user has for the respective category as determined from a preexisting profile of the user, said score further reflecting a degree to which the respective category is relevant to the search query; and

presenting the plurality of categories- to the user, together with associated responsive items, in an order which depends upon the category scores.

19. The method of claim 18, wherein the preexisting profile reflects item purchases made by the user.

20. The method of claim 18, wherein the category score for a given category depends upon item popularity scores of responsive items within the given category.

21. The method of claim 18, wherein the category score for a given category is based further upon a number of responsive items falling within the given category relative to a total number of items falling within the given category.

22. A web page generated according to the method of claim 18 represented within a computer, said web page displaying responsive items arranged by category in said order.

23. A server system that processes search queries according to the method of claim 18.

24. A computer readable medium which stores a computer program that embodies the method of claim 18.

25. A method of assisting users in locating items that are arranged by category within an electronic catalog, the method comprising:

monitoring actions performed by a plurality of users with respect to items in the electronic catalog to generate item usage data;

calculating popularity scores for items in the catalog using said item usage data;

receiving a search query specified by a user, and identifying a set of items in the electronic catalog that are responsive to the search query ("responsive items"), including items from multiple categories;

selecting an order in which to present the multiple categories to the user such that the order is dependent upon the popularity scores of responsive items, said order reflecting predicted levels of relevance of the categories to the search query; and

generating a search results display in which the multiple categories are presented in said order together with corresponding responsive items.

26. The method as in claim 25, wherein the order in which the categories are presented is selected based further upon a number of responsive items falling within a given category relative to a total number of items in the given category.

27. The method of claim 25, wherein the order in which the categories are presented is based further upon a profile of the user such that the order is personalized for the user.

28. A search results display generated according to the method of claim 25 and represented within a computer.

29. A server system that embodies the method of claim 25.

* * * * *